

THE EXPERT'S VOICE® IN OPEN SOURCE

Based on the  
latest GIMP 2.4  
release

# Beginning GIMP

From Novice to Professional

*Create, manipulate, and convert professional-quality images  
with the powerful GNU Image Manipulation Program.*

Akkana Peck

apress®

# Beginning GIMP

From Novice to Professional



Akkana Peck



## **Beginning GIMP: From Novice to Professional**

**Copyright © 2006 by Akkana Peck**

All rights reserved. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage or retrieval system, without the prior written permission of the copyright owner and the publisher.

ISBN-13 (pbk): 978-1-59059-587-9

ISBN-10 (pbk): 1-59059-587-4

Printed and bound in China 9 8 7 6 5 4 3 2 1

Trademarked names may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, we use the names only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

Lead Editor: Matt Wade

Technical Reviewer: Guillermo S. Romero

Editorial Board: Steve Anglin, Dan Appleman, Ewan Buckingham, Gary Cornell, Jason Gilmore, Jonathan Hassell, James Huddleston, Chris Mills, Matthew Moodie, Dominic Shakeshaft, Jim Sumser, Matt Wade

Project Manager and Production Director: Grace Wong

Copy Edit Manager: Nicole LeClerc

Copy Editor: Chandani Thapa

Assistant Production Director: Kari Brooks-Copony

Production Editor: Kelly Winkvist

Compositor: Kinetic Publishing Services, LLC

Proofreader: Liz Berry

Indexer: Broccoli Information Management

Artist: Kinetic Publishing Services, LLC

Cover Designer: Kurt Krames

Manufacturing Director: Tom Debolski

Distributed to the book trade worldwide by Springer-Verlag New York, Inc., 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax 201-348-4505, e-mail [orders-ny@springer-sbm.com](mailto:orders-ny@springer-sbm.com), or visit <http://www.springeronline.com>.

For information on translations, please contact Apress directly at 2560 Ninth Street, Suite 219, Berkeley, CA 94710. Phone 510-549-5930, fax 510-549-5939, e-mail [info@apress.com](mailto:info@apress.com), or visit <http://www.apress.com>.

The information in this book is distributed on an “as is” basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author(s) nor Apress shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this work.

The source code for this book is available to readers at <http://www.apress.com> in the Source Code section. You will need to answer questions pertaining to this book in order to successfully download the code.

*To Dave.*



# Contents

About the Author . . . . .	xv
About the Technical Reviewer . . . . .	xvi
Acknowledgments . . . . .	xvii
Introduction . . . . .	xix
<b>CHAPTER 1 Get to Know the GIMP . . . . .</b>	<b>1</b>
What Is the GIMP? . . . . .	1
A Tour of the Important GIMP Windows and Menus . . . . .	2
The Toolbox Window . . . . .	2
The Layers Dialog . . . . .	3
The Tool Options Dialog . . . . .	4
Image Windows . . . . .	5
Grouping Dialogs by Docking . . . . .	10
GIMP Settings You Can Customize . . . . .	13
Environment . . . . .	13
User Interface and Keyboard Shortcuts . . . . .	14
Theme . . . . .	15
Help System . . . . .	16
Tool Options . . . . .	16
Toolbox . . . . .	18
Default Image . . . . .	18
Default Grid . . . . .	20
Image Windows . . . . .	20
Display . . . . .	21
Color Management . . . . .	21
Input Devices . . . . .	22
Window Management . . . . .	22
Folders . . . . .	23
A First GIMP Project . . . . .	23
Summary . . . . .	26

<b>CHAPTER 2</b>	<b>Improving Digital Photos</b> .....	27
	Opening Files .....	27
	Rescaling .....	29
	What Size Should You Choose for Your Images? .....	30
	The Scale Tool .....	32
	Saving Files .....	32
	Image File Types .....	33
	JPEG .....	34
	GIF .....	34
	PNG .....	35
	XCF .....	35
	TIFF .....	35
	Other Formats .....	36
	Experimenting with JPEG and GIF Settings .....	36
	JPEG Quality Settings .....	37
	GIF and Indexed PNG Quality Settings .....	39
	Cropping .....	42
	Brightening and Darkening .....	44
	Brightness-Contrast .....	44
	Levels .....	47
	Curves .....	52
	Other Exposure Adjustments .....	54
	Rotating .....	58
	Rotating by Multiples of 90 Degrees .....	58
	Free Rotation .....	58
	Sharpening .....	62
	Fixing Red-eye .....	66
	Summary .....	71
<b>CHAPTER 3</b>	<b>Introduction to Layers</b> .....	73
	What Is a Layer? .....	74
	Using Layers to Add Text .....	76
	The Text Tool .....	78
	Using the Move Tool .....	82
	Changing Colors .....	85
	The GIMP's Color Chooser .....	88
	Simple Effects Using Layers .....	89
	Drop Shadows .....	89
	Linking Layers Together .....	93

Performing Operations on Layers . . . . . 94

Using Layers for Copy and Paste . . . . . 96

    Select Only the Part You Want to Paste . . . . . 96

    Paste into the Image . . . . . 98

    Scale the Pasted Layer . . . . . 99

A Tour of the Layers Dialog . . . . . 101

    Title Area . . . . . 102

    Tabs . . . . . 102

    Layer Mode . . . . . 102

    Opacity . . . . . 102

    Layers List . . . . . 103

    Layer Buttons . . . . . 104

    Layer Context Menus . . . . . 105

Bonus Project: Making Simple GIF Animations . . . . . 107

    Create a Base Image . . . . . 107

    Add the Animation Frames . . . . . 108

    Give Each Frame Its Own Background . . . . . 110

    Test Your Animation . . . . . 113

    Animation for the Web: Save as GIF . . . . . 113

Summary . . . . . 115

**CHAPTER 4 Drawing . . . . . 117**

A New Image . . . . . 117

Using Layers for Drawing . . . . . 119

Drawing Lines and Curves . . . . . 121

    Drawing Hard Edged Lines: The Pencil Tool . . . . . 121

    Brushes . . . . . 122

    Drawing Tool Options . . . . . 127

    Drawing Fuzzy or Smooth Lines: The Paintbrush Tool . . . . . 128

    The Airbrush Tool . . . . . 130

    The Ink Pen Tool . . . . . 131

The Eraser Tool . . . . . 134

Drawing Rectangles, Circles, and Other Shapes . . . . . 135

    Defining Regions . . . . . 135

    Outlining Selections: Stroking . . . . . 136

    Free Select with the Lasso . . . . . 138

    Filling Regions . . . . . 138

    The Bucket Fill Tool . . . . . 139

    Patterns . . . . . 142

    Gradients . . . . . 143

A Drawing Project . . . . .	147
Drawing a Tree . . . . .	147
Making a Box Using Perspective Transformation . . . . .	153
Preparing the Planter Box . . . . .	161
Increasing Canvas Size . . . . .	161
Plant the Tree . . . . .	164
Final Touch-Ups . . . . .	164
Summary . . . . .	166
<b>CHAPTER 5 Selection . . . . .</b>	<b>167</b>
Working with Selections . . . . .	168
Marching Ants . . . . .	168
The Select Menu . . . . .	168
Moving Selections . . . . .	170
Select by Color and Fuzzy Select . . . . .	171
Select Contiguous Regions . . . . .	174
Bezier Paths . . . . .	174
Defining a Path . . . . .	175
The Paths Dialog . . . . .	177
Curved Paths . . . . .	184
Adding Nodes or Segments and Moving Paths . . . . .	185
Moving or Modifying an Existing Path . . . . .	185
The Intelligent Scissors . . . . .	186
Modifying Selections with Selection Modes . . . . .	187
The QuickMask . . . . .	191
Highlighting Foreground Objects . . . . .	196
Using Channel Masks to Save a Selection . . . . .	197
Layer Masks . . . . .	203
Extracting Foreground Objects with SIOX . . . . .	213
Summary . . . . .	217
<b>CHAPTER 6 Erasing and Touching Up . . . . .</b>	<b>219</b>
Dimming Highlights with Dodge and Burn . . . . .	219
Dodging . . . . .	221
Burning . . . . .	222
Smudging Blemishes Away . . . . .	225
The Clone Tool, For More Difficult Jobs . . . . .	229
Setting the Clone Source . . . . .	232
Fine-tuning a Clone Job . . . . .	235
Clone Tool Options . . . . .	236

Copying Small Regions . . . . . 237

Sharpening or Blurring Specific Regions: The Convolve Tool . . . . . 240

    Sharpening with the Convolve Tool . . . . . 242

    Blurring with the Convolve Tool . . . . . 242

Blurring Backgrounds with Gaussian Blur . . . . . 245

Correcting Color Balance . . . . . 247

    Hue-Saturation . . . . . 248

    Color Balance . . . . . 250

    Using Curves or Levels for Balancing Colors . . . . . 251

Summary . . . . . 253

**CHAPTER 7 Filters and Effects . . . . . 255**

Image Window Filters vs. Toolbox Xtns . . . . . 255

Filters for Images . . . . . 256

    Tools vs. Plug-Ins . . . . . 256

    The Filters Menu . . . . . 257

    Blur . . . . . 257

    Enhance . . . . . 259

    Distorts . . . . . 259

    Light and Shadow . . . . . 266

    Noise Filters . . . . . 268

    Edge-Detection Filters . . . . . 270

    “Generic” Filters . . . . . 272

    Combine . . . . . 272

    The Artistic Filters . . . . . 273

    The Map Filters . . . . . 279

    Adding Patterns to a Layer . . . . . 285

    Filters to Help Make Web Pages . . . . . 290

    Animation Helpers . . . . . 292

    Alpha to Logo . . . . . 292

    Decor . . . . . 293

Scripts to Make New Images: The Xtns Menu . . . . . 294

Summary . . . . . 295

**CHAPTER 8 Color Manipulation, Channels, and Layer Modes . . . . . 297**

A Review of RGB and CMY Color . . . . . 298

    Additive Colors . . . . . 298

    Subtractive Colors . . . . . 299

    The Relationship Between Additive and Subtractive Systems . . . . . 299

    Color Depth . . . . . 300



The Indexed Palette .....	301
Representing RGB Color in Web Pages .....	301
Working in HSV .....	302
Why Use HSV? .....	303
Working for Print: CMYK .....	303
CMYK for Professional Printers .....	304
GIMP's Alternative Color Choosers .....	304
The Triangle Color Selector .....	305
The CMYK Tab .....	305
The Watercolor Selector .....	306
Working with Grayscale or Black and White .....	306
Methods of Measuring Brightness .....	307
Grayscale Mode .....	308
Desaturate .....	308
Hue-Saturation .....	308
Decompose .....	309
Channel Mixer .....	311
Coloring Monochrome Images and Making Sepia Photos .....	311
Automatic Conversion with the "Old Photo" Filter .....	312
Manual Conversion for Fine Control .....	313
Using Threshold to Clean Up Scans .....	315
Indexed Color .....	317
A Typical Indexing Problem: Choosing the Palette .....	318
Dithering an Indexed Image .....	319
Redesigning for Better Indexed Results .....	320
Pick Colors from the Image .....	322
The Color Channels .....	323
How to Use Color Decomposition .....	324
Use Threshold and QuickMask to Improve the Selection .....	326
Decomposing to HSV .....	327
Color Profiles .....	330
Summary .....	331

## CHAPTER 9 **Advanced Drawing** .....

Useful Mask Tricks .....	333
Making Text "Fade Out" .....	333
Making a Fuzzy Border .....	334
Even Bigger Fuzzy Borders .....	337
Interpreting the Mask Border Color .....	340
Layer Modes .....	341

A Quick Tour Through All of the GIMP's Layer Modes . . . . .	343
Addition, Subtract, and Difference . . . . .	344
Multiply and Divide . . . . .	345
Dodge and Burn, Screen and Overlay . . . . .	345
Hard and Soft Lights . . . . .	346
Darken or Lighten Only . . . . .	347
Grain Extract and Grain Merge . . . . .	347
Hue, Color, Saturation, and Value . . . . .	347
Creating Depth: Drawing with Layer Modes . . . . .	348
Combining Layer Modes: Making 3-D Letters . . . . .	351
Drawing Realistic Shadows . . . . .	357
Using an Object to Cast Its Own Shadow . . . . .	358
Transparency: Add the Final Tweak . . . . .	359
Realism and Multipoint Perspective . . . . .	360
Single-Point Perspective . . . . .	361
Two-Point Perspective . . . . .	361
Adding Reflections and Shading . . . . .	363
Making Brushes, Patterns, and Gradients . . . . .	368
Making Brushes . . . . .	368
Making Patterns . . . . .	374
Making Gradients . . . . .	377
Summary . . . . .	379
<b>CHAPTER 10 Advanced Compositing . . . . .</b>	<b>381</b>
Colorizing Images . . . . .	381
Improving Photos by Self-Compositing with Layer Modes . . . . .	384
Using Screen Mode for Dark Images . . . . .	385
Using Overlay or Hard Light When Light is Flat . . . . .	386
Using Overlay or Burn to Cut Through Haze . . . . .	387
Making Photos into Art: Self-Compositing with Modifications . . . . .	390
Making “Drawings” and Other Effects Using Layer Offset . . . . .	390
Adding Blurs and Other Tricks for Artistic Effects . . . . .	395
Compositing Unrelated Images . . . . .	396
Using Soft Light for Combining Images . . . . .	396
Using Overlay for Dark Images . . . . .	397
Using Screen to Get a Lighter Effect . . . . .	399
Using Addition to Complement Light and Dark . . . . .	399
Using Subtract to Make a Cut-out Mask . . . . .	400
Making Eerie Colors with Burn . . . . .	402
Using Grain Merge to Add Texture . . . . .	403

Stacking Images .....	404
Reducing Noise .....	405
Loading All the Images as Layers .....	405
Registering the Images Using Difference Mode .....	405
Increasing Light by Additive Stacking .....	406
Increasing Contrast by Multiplicative Stacking .....	407
Increasing Resolution by Averaging .....	409
Stitching Panoramas .....	409
Shooting the Images .....	410
Decide on a Resolution .....	411
Calculate Your Expected New Image Size .....	411
Load the First Two Images .....	412
Make a Gradient Layer Mask .....	412
Fine-Tune the Position of the Second Image .....	413
Adjust the Layer Mask .....	414
Adjust with Other Tools if Necessary .....	415
Add the Rest of the Images .....	415
Final Adjustments .....	415
Summary .....	416
<b>CHAPTER 11 Plug-Ins and Scripting .....</b>	<b>417</b>
Plug-Ins .....	418
The Plug-In Browser .....	418
Finding Plug-Ins .....	419
Installing a Plug-In .....	420
Writing GIMP Scripts .....	422
Script-Fu (= Scheme = Lisp) .....	422
Python .....	422
Perl .....	423
Reading and Modifying a Script .....	423
User Interface Options for Scripts .....	437
A Python Script .....	439
Finding Out How to Do Things: The GIMP PDB .....	442
A Perl Script .....	443
Writing a C Plug-In .....	444
A Tour Through Zealous Crop .....	445
The run Procedure .....	447
The Routine That Does the Work .....	449
Getting Pixels: Tiles and Pixel Regions .....	450
Searching the Image for Blank Rows and Columns .....	452

Writing to Pixel Regions .....	453
Making a New Plug-In .....	454
Layer Coordinates vs. Image Coordinates .....	455
Setting Boundaries .....	456
Handling the Edges One at a Time .....	457
Summary .....	458
<b>CHAPTER 12 Additional Topics .....</b>	<b>459</b>
Printing with Gimp-Print or Gutenprint .....	459
Setting Your Printer Model .....	460
Set Print Size, Paper Type, and Resolution .....	461
Setting the Image's Position, Orientation, and Size .....	462
Making the Print: The Buttons at the Bottom .....	463
Fine-Tuning: Image/Output Settings .....	463
Print to File .....	464
Screen Shots and Scanners: the Acquire Menu .....	465
Scanning from GIMP .....	465
Screen Shots .....	465
GIMP Configuration Files .....	467
gimprc .....	468
Changing Preview Sizes with gtkrc .....	471
GIMP's Window Positions and Configuration: sessionrc .....	472
Menu Entries and Key Bindings: menurc .....	473
Other Files .....	473
Folders Inside the GIMP Profile .....	474
Additional Resources .....	475
Official GIMP Documentation .....	475
The Wiki .....	477
Other Tutorial Sites .....	477
Mailing Lists .....	478
Reporting Bugs .....	479
Finding Source Code at cvs.gnome.org .....	481
Where to Find Freely Available Images .....	481
Art Sites and Imaging Contests .....	482
Summary and Conclusion .....	483
<b>APPENDIX A Getting and Installing the GIMP .....</b>	<b>485</b>
Requirements .....	485
Where to Find the GIMP .....	486
Notes on the OS-specific Installation Appendixes .....	486

<b>APPENDIX B</b>	<b>Installing the GIMP on Windows</b> .....	487
	The Quick Start Version .....	487
	The Longer Version: Getting the GIMP .....	488
	Installing GTK+ .....	489
	Installing the GIMP .....	489
	Cleaning Up .....	491
<b>APPENDIX C</b>	<b>Installing the GIMP on Mac</b> .....	493
	Requirements .....	493
	The Quick Start Version .....	493
	The Longer Version: Getting the GIMP .....	494
	If You Need to Get X11 .....	495
	If You Need to Install X11 .....	496
	Installing Gimp.app .....	496
	Other GIMPs and X11s for OS X .....	497
	Cleaning Up .....	498
<b>APPENDIX D</b>	<b>Installing the GIMP on Linux or Unix</b> .....	499
	Installation Notes for Various Distributions .....	499
	Red Hat .....	500
	SuSE .....	500
	Debian Variants .....	500
	Linspire .....	500
	Mandriva .....	500
	FreeBSD and NetBSD .....	500
	Other Unix Versions .....	501
<b>APPENDIX E</b>	<b>Building from Source</b> .....	503
	Getting the Source .....	503
<b>APPENDIX F</b>	<b>First Startup</b> .....	507
<b>INDEX</b> .....		509

# About the Author




■ **AKKANA PECK** is a freelance software developer who has been working with open source software for over 20 years, and using the GIMP for 8 of those years. She has worked for a diverse collection of research labs and companies, including Mozilla development at Netscape, as well as participating in projects at Silicon Graphics, Sun, HP, Apple, and others.

She has written imaging software ranging from GIMP plug-ins to photo viewing to cartography to scientific visualization, as well as articles and how-tos on Linux, astronomy, and other topics for various publications and websites.

About a decade ago, her longstanding hobby of photography spawned an interest in digital imaging and creating photos for the web. Frustration with the existing Windows tools, combined with a switch to Linux as her primary platform, led to fiddling with the basics of the GIMP. After a long period of resistance she was finally persuaded to go beyond basic photo cropping and resizing and try “that layer thing,” and the rest is history.

Akkana is a longtime member of the GIMP community and an occasional contributor to the GIMP’s source code. She enjoys hiking, mountain biking, and astronomy as well as filling up her hard drive with digital images that she can never seem to bring herself to delete, mostly of scenes from local trails and from travels in the desert southwest. She lives in San Jose, California, with her husband and a motley assortment of old computers.

# About the Technical Reviewer



A native of Madrid, Spain, **Guillermo S. Romero**, while always technically oriented, became interested in art through building and painting scale kits as a child, and that has shaped his life ever since. He graduated from Universidad Politécnica de Madrid, with a degree in Telemática, a specialization of Ingeniería Técnica de Telecomunicaciones. While studying, he began to explore Linux, POV-Ray, GIMP, and Blender, mixing technology and art. This knowledge led to some freelance jobs in the form of articles for magazines and reviews for book publishers.

After a few pure telecommunication jobs, he decided that his career should mix technology and art. Currently working for a small publisher dedicated mostly to photography and scale kit magazines and books, he is discovering that drawing with a pen can be learned, step by step, with the right people around providing guidance. His other hobbies are reading, playing in a b&w photography laboratory, walking, and swimming.

# Acknowledgments

**F**irst of all, thanks to my husband, David North. Not only did he help enormously with formatting, proofreading, finding sample photos, and endless testing of GIMP quirks on multiple platforms, but he also put up with the neuroses of a stressed out first-time author.

Equally important is Guillermo Romero, the book's technical reviewer, without whom this book would not have been possible. He patiently waded through first drafts, flagged my silly errors, asked probing questions, and took the time to teach me a lot about digital art. I'm not an artist yet, but I'm learning.

Thanks to all the GIMP developers for using their talents and spare time to create such a wonderful program, and for spending even more time explaining details of the program to folks who need help. Some of the most active (in order by CVS check-ins): Sven Neumann, Michael Natterer, Bill Skaggs, DindinX, Manish Singh, Karine Delvare, Simon Budig, Tor Lillqvist, Nathan Summers, Michael Schumacher, Kevin Cozens, Maurits Rijk, Seth Burgess, Raphaël Quinet, Helvetix Victorinox, Hans Breuer, Jay Cox, Øyvind Kolås, Jernej Simoncic, and Aaron Voisine. And thanks to Carol Spears, whose online tutorials and explanations cleared up all sorts of details that weren't explained anywhere else.

Thanks to Pat Peck (hi, Mom!) and the members of Linuxchix who were willing to be guinea pigs for the online GIMP course that indirectly led to this book. And to two other Linuxchix: Carla Schroder, author extraordinaire, who reassured me that I could actually take on this task and was a great source of writing tips and witty remarks; and Dana Sibera, for laughs and an appreciation of what a real artist can do with (or to) a photo.

Thanks to Lisa and Evan Avery, for letting me use the photo of their son Ethan; to Cathleen Wang Blythe and the world's cutest dog, "Polarbear"; and to Dave Nakamoto and Dragan Stanojević (Nevidljvi) for helping out with samples of image stacking (which unfortunately didn't make it into the final version due to space considerations).

Last but not least, many thanks to the folks at Apress for helping me through every step of the process, and for putting up with a new author's never-ending tweaks and revisions. In particular, thanks to open source editor Matt Wade, project manager Grace Wong, insightful copy editor Chandani Thapa, attentive production editor Kelly Winquist, and eagle-eyed proofreader Elizabeth Berry.





# Introduction

**S**o you want to learn image editing!

Maybe you've been shooting lots of pictures with your digital camera and want to learn how to make them look great, or prepare them for the web.

Maybe you're interested in creating drawings or cartoons. Or you're already editing images, but you want to get better . . . even learn some of the theory behind imaging.

In any case, you've been hearing about the GIMP—the GNU Image Manipulation Program—and you're ready to learn how to use it and get the most out of it.

This book is ideal for anyone attempting image editing for the first time, but it goes deep enough to satisfy the intermediate GIMPster who knows the basics but needs more. Plus, it contains enough hidden tricks and reference matter to gratify even power users. Anyone with an interest in digital art and a willingness to explore should find lots of useful tips and fun projects throughout the book.

## Structure of the Book

The first few chapters assume no knowledge at all of the GIMP or of any other image editing program. Later chapters will assume you've picked up these basics, and will build on them.

Chapter 1 introduces the GIMP's various windows, menus, and interface conventions. GIMP has a somewhat different user interface from most other programs, so it helps to get an idea of how its windows and dialog boxes work together and how the menus are structured.

Beginning with Chapter 2, you'll dive into practical image editing tasks. Chapters 2 through 6 each cover a different category of image operations—layers, drawing, selection, touching up—and each one describes a series of different and related techniques.

Chapter 7 takes you on a whirlwind tour of the special effects available in the GIMP. Then you'll dive into more advanced topics in Chapters 8 through 10, which cover subjects such as color theory, layer mode effects, shading, perspective, image stacking, and panoramas.

Chapter 11 explores scripts and plug-ins: how they work, how to install existing ones, and how to take one and tweak it to do something slightly different. I hope that even people with no programming experience will at least take a look at the scripting sections; there's a lot you can do with GIMP scripts, even if you've never written a line of code before.

For the exercises throughout most of the book, you will want to have some digital photographs handy. If you don't already have a digital camera or a collection of scanned photographs, “Additional Resources” in Chapter 12 offers a list of places on the web to find all sorts of terrific photos you can use in your own projects. The rest of Chapter 12 covers a few topics that don't fit anywhere else in the book, plus web resources for finding more GIMP information.

If you don't already have the GIMP installed on your computer, begin by skipping ahead to Appendix A for an outline of how to install the GIMP on the most common platforms it supports—Windows, Mac OS X, and Linux. If you ever decide you want to build the GIMP from source, Appendix F has tips on how to do just that.

## A Note on Versions

This book was written based on GIMP 2.3 development builds, slated to become the next stable release, 2.4. However, it's intended to be useful with GIMP versions 2.2 and 2.0 as well.

Some features have changed during the development of 2.3, and some of the menus were restructured to group items in a more logical way. When possible, I mention locations of menu items in the older GIMP 2.0/2.2 menu structure as well as their new location in GIMP 2.4.

A few features, such as the SIOX selection tool, are new for 2.4. Users of 2.2 can skip those sections, or try installing the latest GIMP in order to try them out. Rest assured that even if you don't have those features yet, you will eventually!

## Downloads and Feedback

When learning image editing, it's always helpful to have lots of examples to work with. Some of the images that appear in the book are available in the GIMP's native XCF format on the Apress web site, along with any scripts used to create them. You'll also find source code there for the scripts and plug-ins presented in Chapter 11, along with any (gasp!) errata. Just point your browser to *www.apress.com* and search for this book.

Please email any feedback or suggestions to *akkana@gimpbook.com*.

I hope you enjoy your introduction to the GIMP!



# Get to Know the GIMP

**W**elcome to the GIMP, the premiere open source image-editing program!

This chapter will introduce you to the GIMP's most important windows, menus, and settings, and will offer some tips along the way. I'll cover the following topics:

- What is the GIMP?
- A tour of the important GIMP windows and menus
- Some GIMP settings you can customize, and why you might want to
- A first GIMP project

At the end of this chapter, you'll get a chance to try a simple project: combining two images into one.

## What Is the GIMP?

The *GNU Image Manipulation Program*, more commonly known as the GIMP, is a program for creating and editing bitmapped images. That means it's well suited for editing digital photographs and typical web graphics. You can also use it to make some pretty amazing drawings.

The GIMP is a complex program suitable for professional artwork, but it's also just plain fun—a place to play with pretty pictures and let your imagination run wild. As you read this book, please keep that in mind. Make a point of “fooling around” and having fun with your image-editing projects.

The GIMP arose out of a 1995 computer science class project by Spencer Kimball and Peter Mattis. Within a year, it had grown into a popular open source project, with users and contributors from around the world.

Today, the GIMP is used by countless people worldwide. It runs on most Windows systems, Mac OS X, Linux, and most versions of UNIX. The program is free software in both senses of the word “free”: you don't have to pay anything to use it, and the GIMP's inner workings, or “source code,” are available for anyone to examine, contribute to, distribute, or learn from. The latest is always available from <http://www.gimp.org>.

The GIMP stands as a shining example of the power of free, open source software. It's written and maintained by volunteers who keep in touch by means of internet relay chat (IRC) and mailing lists, despite time zone and language differences. Users of the GIMP who want to

get more involved can participate in the GIMP's mailing lists, contribute bug fixes or new features, design new brushes or patterns, or write their own plug-ins and scripts to share with the world. Or they can just enter goofy photo contests to show what cool things can be done with the GIMP!

---

**Note** The GIMP is constantly evolving. The screenshots in this book are based on an early version of GIMP 2.4, though most of the information here applies to any version of the GIMP from 2.0 on.

---

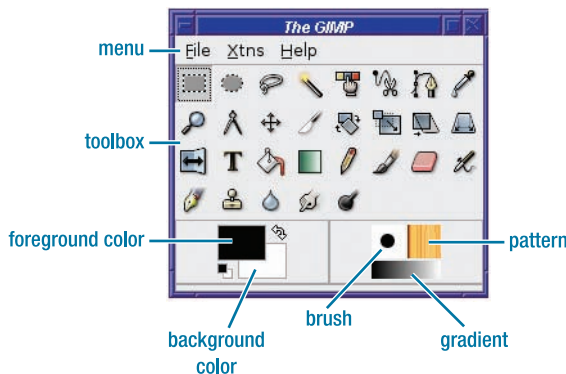
## A Tour of the Important GIMP Windows and Menus

The GIMP normally enables you to keep several windows open at once, rather than having all its functions available from one window. This gives you control over where the various windows are positioned and which ones remain visible.

The following sections describe some of the most important windows GIMP uses: the Toolbox window (and its sibling, the Tool Options dialog), the Layers dialog, and the image windows.

### The Toolbox Window

The Toolbox is the GIMP's main window. No matter how many images you have open at once, you'll only have one Toolbox. In its most basic form, it looks like Figure 1-1.



**Figure 1-1.** *The GIMP Toolbox*

The Toolbox window has four parts: a menu bar, the Toolbox buttons, the color swatches, and buttons for the current brush, pattern, and gradient.

You can resize the Toolbox to any size you like, and the tool buttons will adjust their positions. As new versions of the GIMP are released, the tool buttons sometimes change position, so it's best not to rely on a tool being in any particular place in the Toolbox.

The menu bar offers operations that don't apply to a currently open image. This includes *File* operations such as *New* and *Open*, a *Help* menu for online help, and a menu labeled *Xtns*.

Of particular interest is the *Xtns* menu. This menu gives you access to a menagerie of extensions that can create new images. Inside are submenus offering a collection of scripts that create many different styles of buttons (for use on web pages), text logos, shapes such as a sphere, interesting patterns and brushes, and themes you can use for web pages. Take a few moments to explore some of these functions. You can accomplish quite a bit right here, without knowing any more of the GIMP.

The *Xtns* menu also contains the GIMP's Plug-in Browser. Many of the GIMP's functions are implemented as plug-ins, and it can sometimes be difficult to find just what you're looking for in the menus. The Plug-in Browser lets you search for plug-ins by name, and then tells you where in the menus to look for them. See Chapter 11 for more details.

The Toolbox buttons, occupying the central part of the Toolbox window, represent tools for drawing, changing, and modifying images. The tools are specific to various functions such as drawing, handling text, or moving objects. You'll learn more about them in Chapters 2, 3, and 4. Each tool has a *tooltip* to help you remember its function, which you can see by moving your mouse over each button and pausing for a second or two.

The GIMP's tools are also accessible through the *Tools* menu in the image window, and most tools have an associated keyboard shortcut, which you can learn by looking for the tool in the *Tools* menu. In addition, there's a Tools dialog, accessed via *Dialogs* ► *Tools* from any image window, or *File* ► *Dialogs* ► *Tools* from the Toolbox.

---

**Tip** Not only can you access the GIMP's tools from the Tools dialog, but also you can add or delete tools from the buttons shown in the Toolbox by clicking on the "eyeball" icon next to each tool in the Tools dialog. You can also change the order of the Toolbox buttons by dragging and dropping the tools within the dialog.

---

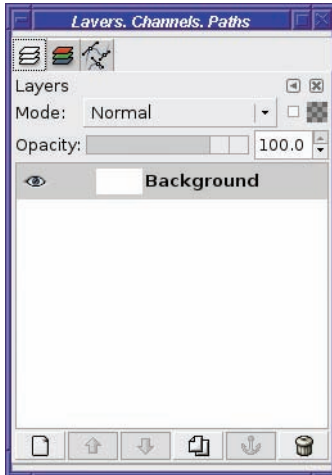
The color swatches allow you to change the foreground and background colors used by the GIMP's drawing and text tools. Initially they're black and white, but clicking on one of the swatches brings up a color chooser that allows you to select any color. This will be explained further in Chapters 3 and 4.

The brush, pattern, and gradient are used for drawing; these will be discussed in more detail in Chapter 4.

## The Layers Dialog

In addition to the Toolbox window, the GIMP brings up a dialog window labeled "Layers, Channels, Paths" (Figure 1-2). For simplicity, I'll refer to this as the Layers dialog.

Your first instinct upon seeing the Layers dialog may be to wonder what it is and close it. Don't! You'll be using layers quite soon (Chapter 3), and they're incredibly useful for all sorts of image operations. If you do close the Layers dialog and want to get it back, you can find it in the *File* ► *Dialogs* submenu in the Toolbox window, along with a collection of other dialogs the GIMP offers.



**Figure 1-2.** *The Layers dialog*

---

**Note** The GIMP has quite a few dialogs in addition to the Layers dialog, and these can be grouped together. You can see a list of them in the *File* ► *Dialogs* submenu of the Toolbox window. See the section “Grouping Dialogs by Docking” later in this chapter.

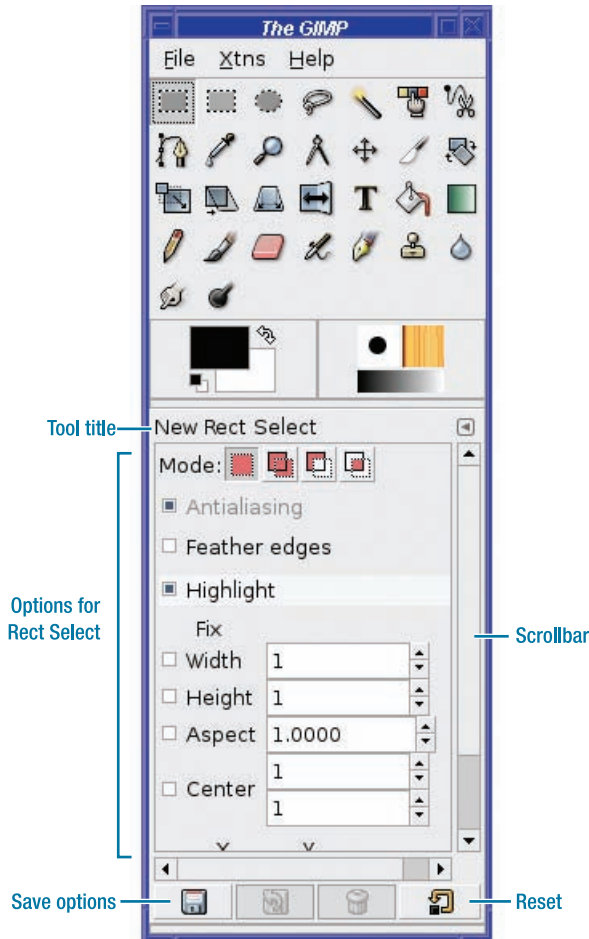
---

## The Tool Options Dialog

Each tool in the Toolbox has options associated with it. These options are collected in a dialog called Tool Options.

The Tool Options dialog is normally docked just below the Toolbox (Figure 1-3), although you can split it off into a separate dialog if you prefer. (See later in this chapter for an explanation of docking under “Grouping Dialogs by Docking.”) The dialog isn’t titled “Tool Options”; instead, the title reflects the current tool, in this case, “Rect Select.” When you choose a new tool in the Toolbox, the contents of the Tool Options dialog immediately change to reflect options that apply to the new tool. If you prefer to use Tool Options as a separate dialog, double-clicking any tool in the Toolbox will bring up its Tool Options dialog.

Some tools have quite a few options. If there are too many options to fit in the window size you’ve chosen, a scrollbar will appear. The most common options are usually listed first, so most of the time you won’t need to scroll to see the options you need.



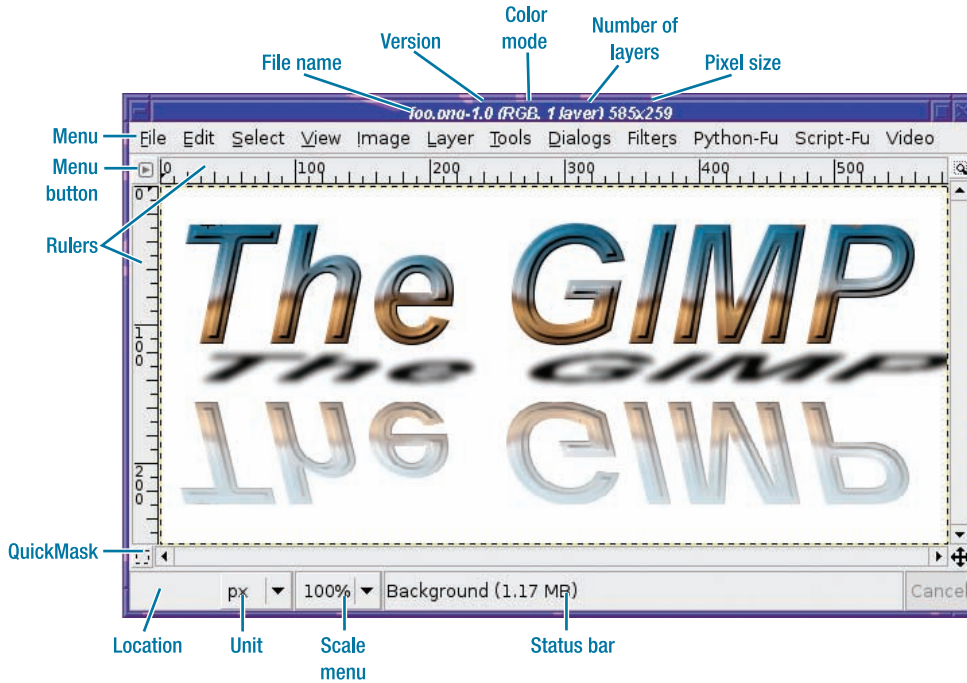
**Figure 1-3.** *Tool Options for Rect Select*

## Image Windows

An image-editing program isn't any fun without images. In the GIMP, every image has its own window (Figure 1-4).

Let's start at the top. The image window's title bar gives information about the image: its file name and internal GIMP identification number, whether or not it's in full color (denoted RGB, for *red, green, and blue*), the number of layers, and the current size of the image. You can change the information in the GIMP's title bar; see the section "Image Windows: Title and Status" later in this chapter for more information.





**Figure 1-4.** An image window

The window also has a menu bar, which offers a selection of operations for working on the image. You can hide the menu bar to offer more space for the image (via Preferences); in that case, use the menu button at the upper-left corner of the window to show the menu, or click the right mouse button on the image.

The rulers at the top and left side of the image window show the position of your mouse in the image. You can also use the rulers to set up *guides*, or grid lines on the image to help you draw exactly where you want (you'll meet guides in Chapter 4).

The *QuickMask* button helps with defining selections. See Chapter 5 for a detailed discussion of how to use the quick mask.

The location area in the lower-left corner gives a precise readout of the current mouse position if your mouse is in the window (otherwise it's blank). The menu button next to it lets you show the mouse position in any unit you find most comfortable, such as pixels, inches, or points.

The *scale menu* shows whether the image is currently being displayed at actual size (100%), shrunk to a smaller size, or magnified to a larger size. You can change the magnification with this menu, via the *View* ► *Zoom* menu, or by using the + and – keys. If you have a mouse with a scroll wheel, you can also zoom by using the Control key in combination with the mouse wheel.

The status bar shows more information about the image: specifically, which layer is selected and how much memory the image is using. The *Cancel* button next to the status bar becomes active when the GIMP is in the midst of a time-consuming operation.

The navigation control in the lower-right corner lets you *pan* around the image instead of using the scrollbars. For large images that don't fit on the screen, panning can be quicker. You can also pan by dragging in the image window while holding down the middle mouse button.

The zoom control toggles whether the image window changes its size when you zoom in or out; you can also change this behavior with a preference.

## Context Menus and Tear-offs

Quite a few operations in the GIMP can be accessed via *context menus*, by using the right mouse button (or Command-click on the Mac if you still have a one-button mouse). For instance, layers in the Layers dialog have a context menu associated with them. Right-click on the line corresponding to a specific layer to see a menu of operations you can run on that layer.

Some menus, such as the menus in image windows, are also available as menu bars at the top of the window. But it's useful to know about the context menus as well, and they offer some features that the menu bars don't. The most useful of these features is *tear-offs* (Figure 1-5).

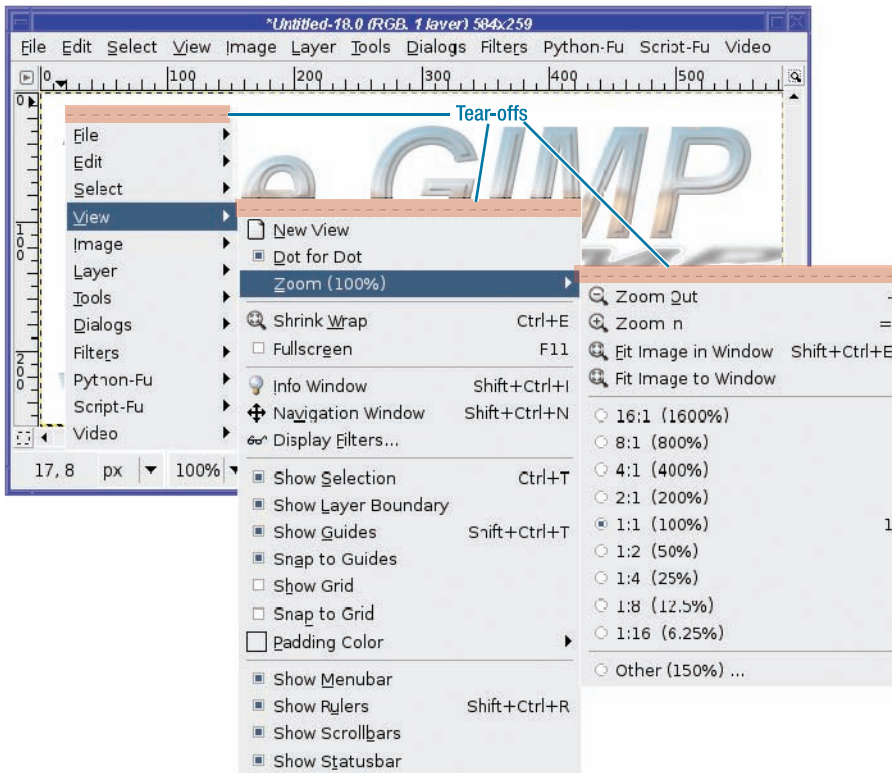
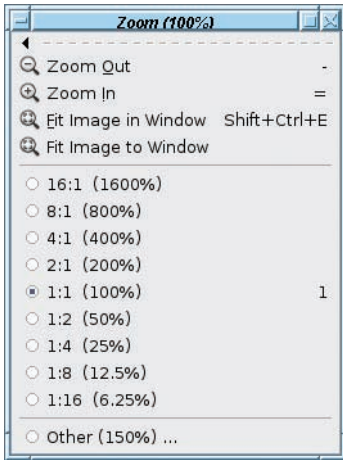


Figure 1-5. Context menus

Figure 1-5 shows a context menu expanded several levels deep. Notice the dotted lines at the top of each menu. These are tear-off lines—think of them as perforations for tearing off the menu underneath. Click on a tear-off, and the menu under it detaches and becomes a separate window. This is a useful way to access functions that you need to run often, without having to navigate the menu system each time. Figure 1-6 shows what a torn-off menu window looks like. You can move this window to a convenient place on the screen for one-click access to the functions inside. Clicking on the dotted tear-off at the top of the window will close the window.



**Figure 1-6.** A torn-off menu

## Image Window Menus

The image window's menus, when combined with the tools in the Toolbox, make up the heart of the GIMP. Everything you'll need to modify an image should be accessible here.

The *File* menu contains familiar file operations: *New*, *Open*, *Save*, *Save As...* Don't neglect the *Open Recent* submenu; you can use it to get back quickly to images you've edited in the past few days.

Edit includes *Cut*, *Copy*, and *Paste* operations, which you will use frequently to move information between images and layers. Other useful items in this menu are: *Clear* (which removes the contents of the current selection) and three variants of *Fill*. Most of the items in the Edit menu have keyboard equivalents. You may find you'll want to learn the shortcuts for these items since they're so frequently useful.

The *Select* menu controls actions that modify the current selection. You can select *All* or *None*, *Invert* the selection (select everything not currently selected, and vice versa), make the selection smaller (*Shrink*) or larger (*Grow*), or its edges fuzzier (*Feather*) or sharper (*Sharpen*).

The *View* menu controls how you see the image. Items in this menu don't change the image that will be saved to disk, merely the way you view it on screen as you edit it. The menu includes a *Zoom* control that lets you zoom in or out on the image (this does not make the

image any bigger or smaller when you save it, it merely changes your view of it), as well as toggles to select whether various attributes such as the selection, guides, rulers, layer boundaries, grids, and the menu bar are visible.

---

**Tip** An easy item to overlook in the *View* menu is the very first one: *New View*. This lets you have more than one view of the same image. For example, you might have one window showing the image zoomed in so you can change individual pixels, and another showing the image at normal size to see what the image will look like when you're finished. They're both still the same image as far as the GIMP is concerned: any changes you make in one view will be reflected in the other.

---

The *Image* menu contains functions that apply to the current image as a whole:

- *Image* ► *Duplicate* makes a new copy of the image, which will appear in a new window: this is not a separate view of the same image, as *New View* would be, but a new image, with a new name. Changes you make to this image will not be reflected in the original image.
- *Image* ► *Transform* provides a quick way to flip an image horizontally or vertically, or to rotate it 90 or 180 degrees. For more complex rotations, use the Rotation tool. You'll learn more about rotation in Chapter 2.
- The rest of the *Image* menu concerns the image's size in pixels or your selected units (*Canvas Size...*), the image's print resolution (*Print Size...*), rescaling of images larger or smaller (this actually changes the image, unlike *Zoom*), some cropping and layer utilities, and utilities to handle guides and the grid. All these functions will be covered thoroughly in subsequent chapters.

The *Layer* menu contains functions that operate only on the current layer. Some of these items look the same as items in the *Image* menu: for instance, *Transform*, *Scale*, and *Autocrop* appear in both menus. But they're different: items in the *Layer* menu only affect one layer and not the whole image. This distinction will become much clearer in Chapter 3, when you start to use layers.

The *Colors* menu contains functions related to the colors of the image, which you will learn about in Chapter 2. Earlier GIMP versions located this menu inside the *Layer* menu.

The *Tools* menu gives you another way to access the tools in the Toolbox window. Most of the time you'll probably keep the Toolbox open and access tools that way, so you won't need this menu.

The *Dialogs* menu gives you access to any dialog, whether or not it's currently visible; it's the same as the *Dialogs* menu in the Toolbox window.

*Filters* provides access to the plethora of image filters and plug-ins available in the GIMP. Most external plug-ins you install will also show up in these menus. (In earlier versions of the GIMP, some of these functions appeared in additional menus such as *Script-Fu* and *Python-Fu*.)

---

**Tip** In the GIMP 2.2 and earlier, it wasn't always easy to predict in which of these menus a particular function would show up. It depended on the language in which the function was written: C, Python, Perl, or Script-Fu. You'll learn about these differences in Chapter 11; for now, either explore on your own, or try using the Plug-in Browser in the Toolbox *Xtns* menu to locate the function you need.

Functions in all of these menus operate on the current image. Chapter 7 will explore them in much greater detail. It's well worth spending a few minutes (which may turn into hours) trying the various filters on an image to see what they do. You can take a photograph and turn it into an oil painting on canvas, a cylinder or cone, a rapidly moving blur, or a line-art sketch. You can add borders, drop shadows, or coffee stains. The possibilities are nearly endless.

---

## Grouping Dialogs by Docking

As mentioned earlier, GIMP dialogs can be grouped together using a process called *docking*. For instance, the “Layers, Channels, and Paths” is really three dialogs, but GIMP normally groups them together as tabs in one dialog.

Docking lets you combine windows in a way that fits your own personal style: you can have lots of little windows, or you can dock windows together to make one large control window. (Image windows will still be separate.)

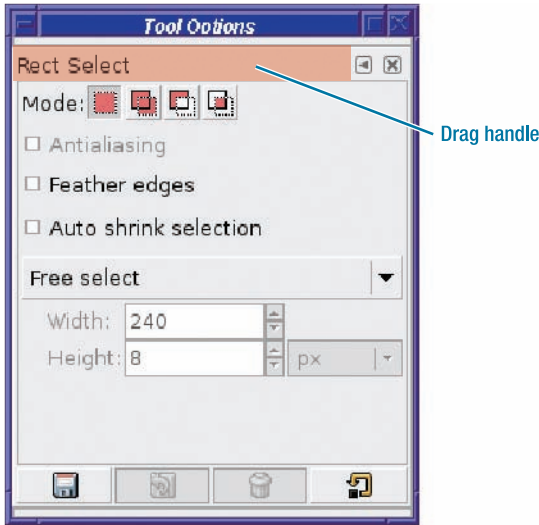
Each GIMP dialog window has a set of *docking bars*, thin horizontal strips you can use to add other dialogs to the window (Figure 1-7). There's usually a docking bar at the top of the window, another one at the bottom, and, if more than one function is docked into the window, an additional docking bar below each function.



Docking bar

**Figure 1-7.** Docking bar in the Toolbox window

Separate dialogs (those not currently docked into another window) also have a drag-handle area (Figure 1-8), including the dock's title and an area to the right of it, which offers a place to drag.



**Figure 1-8.** Drag handle in the Tool Options dialog

To dock a dialog, drag with the left mouse button from the drag handle into the docking bar where you want the dialog to end up. The docking bar will light up in a different color when you drag over it, letting you know that you can release the mouse button and dock your dialog there.

---

**Note** In some Windows versions, you may see a cursor shift rather than a color change. Other than that, it works the same way.

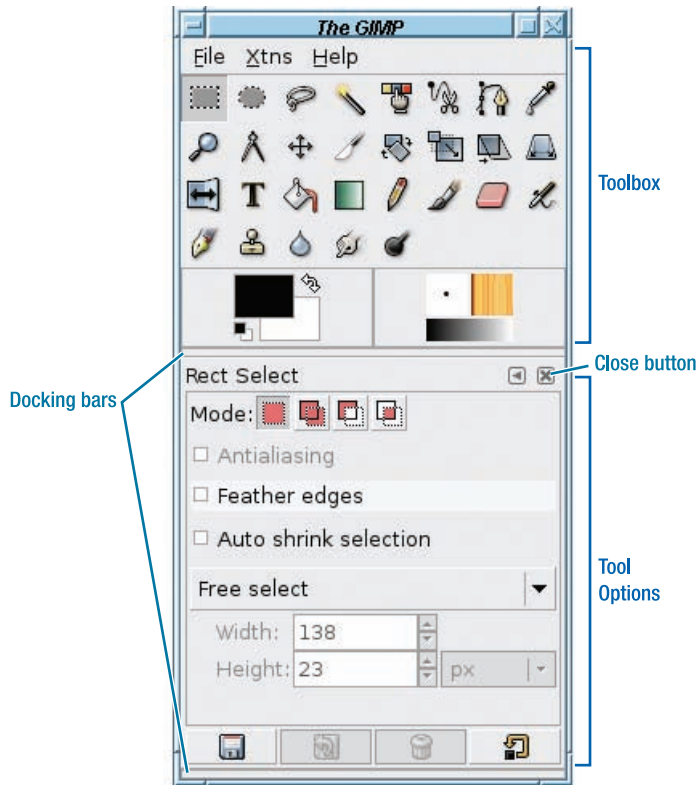
---

Figure 1-9 shows the Toolbox with the Tool Options dialog docked underneath it. This is a very common and popular combination, which will probably be the default when you first start the GIMP.

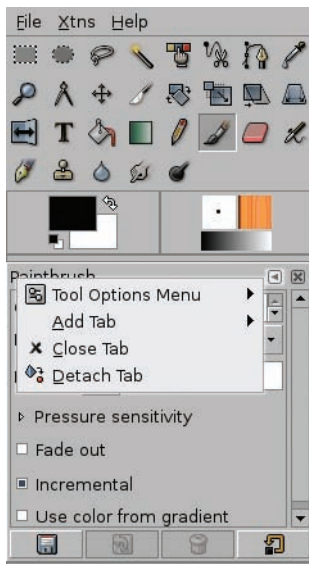
Notice that when a dialog is docked into another window, its drag handle is still there. You can use the handle to drag that dialog out of its current window into another window, or onto the desktop. The “X” at the far-right side of the title area acts as a close button; selecting it will close that dialog (in this case, Tool Options) entirely. To get it back, use the *File* ► *Dialogs* sub-menu from the Toolbox.

What about the arrow button to the left of the close button?

That’s the *Tab* menu (Figure 1-10). It allows dialogs to be added as tabs, as in the “Layers, Channels, and Paths” dialog, or existing tabs to be detached or closed. Tabs can also be re-ordered simply by dragging them. In some dialogs, other options are added to the *Tab* menu.



**Figure 1-9.** Tool Options docked below the Toolbox

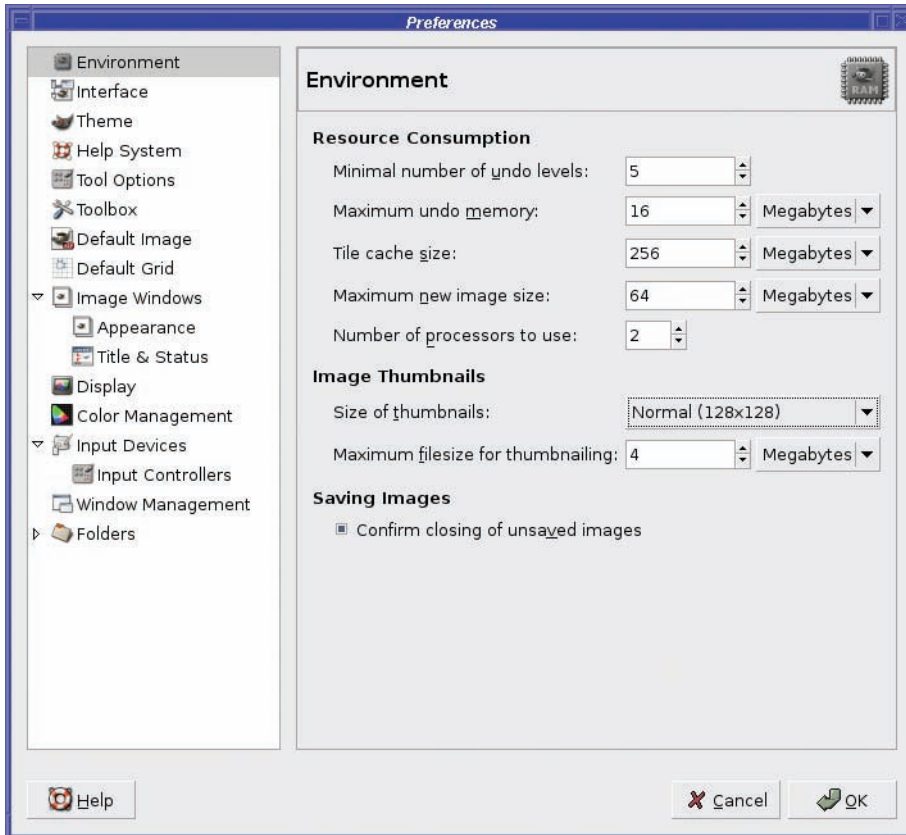


**Figure 1-10.** The Tab menu



## GIMP Settings You Can Customize

Figure 1-11 shows the GIMP's Preferences window (*File* ► *Preferences*).



**Figure 1-11.** *The Preferences window*

On the left side of the Preferences window is a list of categories. Selecting any category in the list shows the preferences associated with that category.

I'll go over each of the categories and some of the useful settings contained within each one.

### Environment

The *Resource Consumption* preferences let you configure several factors affecting the GIMP's memory usage. Most of these features are fairly self-explanatory: increasing them will give you better performance and more flexibility, but will cause the GIMP to consume more RAM. If you work with large images and have plenty of memory, increasing the size of the file cache can improve performance quite a bit, and, of course, having lots of undo levels is always nice if you have memory to spare.

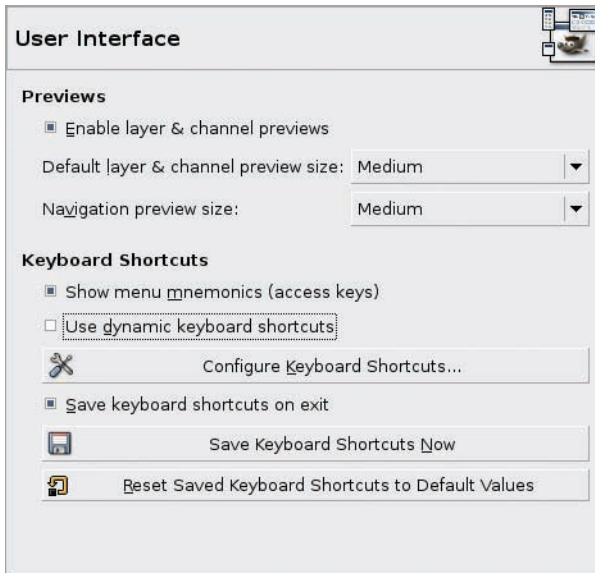


*Image thumbnails* are stored in a common location so that other applications can use the thumbnails GIMP generates. Of course, they do take up some disk space; you can control their size, or disable thumbnail generation entirely, with the thumbnail size preference.

Finally, the GIMP normally asks about unsaved images, if there are any still open when you quit the application. You can disable this behavior here, which will cause any changes to be discarded without prompting. Be careful! Most people will want to leave this option enabled. You can double-click on an image in the dialog to call up that image's window, and you can leave the dialog up as you check the open images to make sure you haven't forgotten to save your work.

## User Interface and Keyboard Shortcuts

The *User Interface* preference category (Figure 1-12; the preference panel is headed *User Interface* though the category on the left says merely *Interface*) lets you configure the way the GIMP's windows appear and behave.



**Figure 1-12.** *Interface preferences*

The first group of options controls previews that will be shown in the Layers and Channels dialogs. If you find that you want a larger or smaller preview size, you can change them here.

The second group concerns keyboard shortcuts. This is one of the GIMP's most useful hidden features. Like most programs, the GIMP has many functions bound to particular shortcut keys. Pressing Control+S in an image window saves the image to a file; pressing Control+N brings up a New Image window; and pressing Control+Q in any window quits the application. When you look at an option in a menu, any keystrokes that will activate that item are shown on the right side of the menu.

However, the GIMP also lets you set up your own custom keyboard shortcuts. If you run a particular function often, rather than navigate the menus each time you use that function, you can bind it to a key. You just have to select the toggle button next to the *Use dynamic keyboard shortcuts* option in the interface preferences.

Once the *Use dynamic keyboard shortcuts* preference is enabled, you can add a shortcut for an action by bringing up the menu containing that action, moving your mouse over the action (without clicking), and pressing the key you want assigned to that action. To delete a shortcut, use Backspace.

For example, suppose you often need to bring up the Preferences window. You can speed that up by triggering it with Alt+P. Click on *File* in the Toolbox window. Move your mouse over *Preferences* (but don't press the mouse button). Now type **Alt+P**. The *Preferences* menu item will change to include the Alt+P shortcut. Now, whenever you press Alt+P in the Toolbox window, the Preferences window will appear.

---

**Caution** A new shortcut will override any previous function which that key combination previously might have called up. So choose your special shortcuts carefully. If you give a new meaning to Control+S, then you will no longer be able to use Control+S to save a file. Binding items to Alt key sequences may interfere with using the Alt key for menu access; if you have problems binding Alt keys and you don't use them to navigate menus, you may want to turn off the *Show menu mnemonics* option. You can view all the current keyboard shortcuts with the *Configure Keyboard Shortcuts* button in the Interface Preferences window, and you can also set them from there.

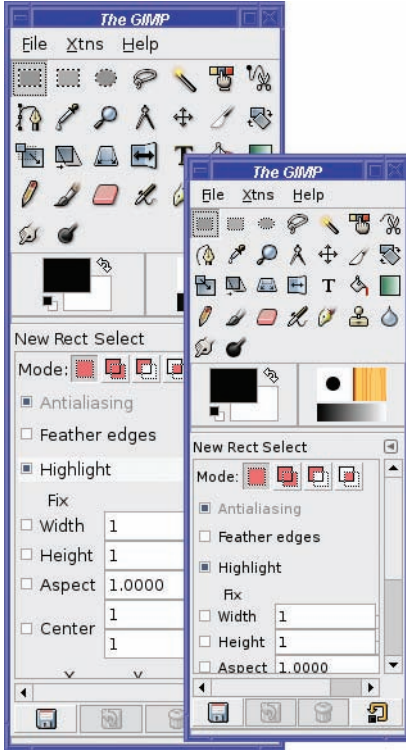
---

## Theme

The *Theme* preference lets you apply a different look to the GIMP application.

The GIMP comes with two themes: *Default* and *Small* (Figure 1-13). The *Small* theme makes the Toolbox and dialog windows smaller, so you have more screen space available to use for image windows.

You can also use this preference to install new themes—for instance, to use a different set of icons for each of the tools, or a less colorful set of icons to distract less from the image being edited. However, there don't seem to be many themes available for the GIMP yet. Try a web search for *GIMP themes* if you want to see what is available.



**Figure 1-13.** *Default (left) and Small themes*

## Help System

The *Help System* preferences let you configure the GIMP's online help. You can view the help in a browser of your choosing, or in a window of the GIMP's making.

---

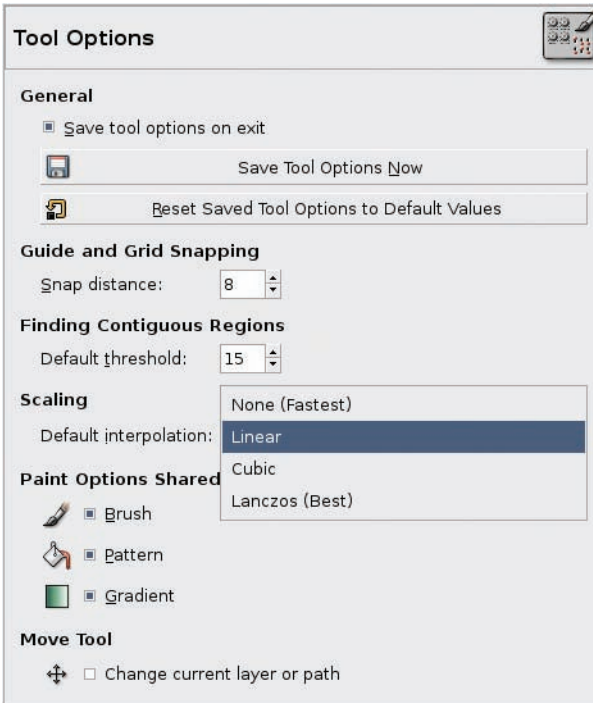
**Caution** Some GIMP installations don't include a Help browser. If you have trouble viewing the GIMP's built-in Help, check the place where you got your GIMP to see if it offers a separate GIMP Help package you can download. Windows and Linux GIMP installations often package the Help files separately, for various technical reasons.

---

The GIMP will also show helpful tips when you start it, unless you deselect *Show tip next time GIMP starts*. You can see these tips at other times by choosing *Help* ► *Tip of the Day* from the Toolbox window.

## Tool Options

*Tool Options* controls some common settings for a few of GIMP's tools (Figure 1-14).



**Figure 1-14.** Tool Options preferences, showing Scaling options

First is a checkbox controlling whether the GIMP saves tool options on exit. With this box checked, any changes you make in Tool Options dialogs will be remembered from one GIMP session to the next. Leave the box unchecked, and the GIMP will reset all the tool options to default values the next time you start up. This is purely a matter of personal preference. If you don't save tool options on exit, there's a button to save the tool options when you get them just the way you want them.

The next option controls how strongly the GIMP will “snap” to the grid (if you have it shown) or guides (if you have any in position). Snapping to the guides or grid can make it easier to put lines, squares, or rectangles in very specific positions.

*Finding Contiguous Regions* controls the default value for the Fuzzy Select tool. See Chapter 5 for a detailed discussion.

The *Scaling* option is worth thinking about right away. When you change the size of an image, it can be scaled larger or smaller in several ways, depending on whether you want the operation to happen more quickly, or whether you want the best quality in the final image. Unless you have a very slow machine, I advise choosing the setting recommended as “Best.” Even on a slow machine, you may still want to use the Best setting, and just be patient when scaling large images; but if it's too slow, try one of the other settings. *Linear* is probably the fastest you'll want to try; *None* will result in lower quality than most people can tolerate.

*Paint Options Shared Between Tools* enables the brush, pattern, and gradient to be set separately for each tool, or tools can all use the same settings. As you become more skillful

with drawing (Chapter 4), you will get a feel for how you want to deal with this issue. The default setting is probably fine for most people.

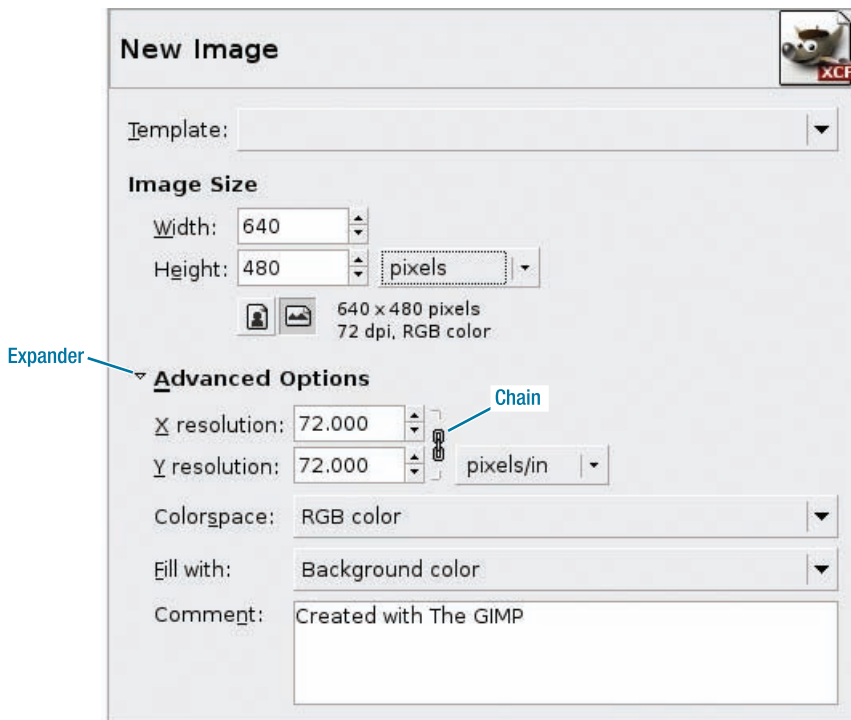
I'll talk about the Move Tool and its options in Chapter 3, since it's used to move layers. For now, leave this setting unchanged.

## Toolbox

The three *Toolbox* options control which items are shown in the Toolbox window under the tool buttons. You can turn off the items you saw at the bottom of the Toolbox window in Figure 1-1, individually or together, and get a smaller Toolbox window consisting only of the menu bar and the collection of tool buttons.

## Default Image

The *Default Image* preferences (Figure 1-15) determine the default look of an image you create with the GIMP's *File* ➤ *New* action.



**Figure 1-15.** *Default Image preferences*

The first item is a list of templates. Most of the templates correspond to common pixel resolutions, such as 640×480, or fixed paper sizes, such as US Letter or CD Cover. You can create custom templates using the Templates dialog, accessed from *File* ➤ *Dialogs* ➤ *Templates* in the Toolbox window.

You can choose a template, then modify it—or simply use your own preferred image size and resolution. The *Image Size* group of options allows you to specify a preferred image's width and height. These can be specified in pixels, or in other units such as inches, millimeters, or points. You can edit the number in the text field directly, or use the up and down arrows next to each text field to make the number larger or smaller. Below the *Width* and *Height* fields are a couple of buttons that show whether the image will be *portrait* (taller than it is wide) or *landscape* (wider than it is tall). Switching between these modes will also swap the width and height values automatically.

The units menu, next to the *Height* field, lets you specify the image's width and height in any unit you prefer. You can work in inches, millimeters, or points if it feels more comfortable.

Since the GIMP is an image-editing program, it always works in pixels, and converts other units into pixels using your current resolution setting. To change the resolution setting, open the *Advanced Options* section if it isn't already visible by clicking on the expander to the left of the words *Advanced Options*.

Resolution is normally specified in pixels per inch (also called dots per inch, or dpi), though you can change that with the units menu. Most often, you will probably want this to be the same as the resolution of your computer monitor (you can get this from the *Display* category in the Preferences window, which is discussed in the section “Display”), but you may want to set a higher resolution for images that are primarily intended for a printer.

Normally, the same resolution is used for both X (horizontal) and Y (vertical) dimensions, so when you change one, the GIMP will change the other; they are “linked together.” However, some printers may use different resolution settings for X and Y, and it might be important for some video and film setups. If you need to unlink two quantities that are chained together in the GIMP, use the “chain link” button. When the link is closed (Figure 1-16), the two quantities are linked together: changing one quantity will cause the other to change. Click on the chain link, and it opens—now the two quantities are independent, and you can change them separately.



**Figure 1-16.** *Chain link icons*

Below the resolution options is a *Colorspace* selector (Figure 1-15). New images can be in full color (RGB) or in grayscale. Most people will want to leave this option as RGB.

*Fill with* lets you control the page color of new images. You can use either the GIMP's current background color, its foreground color, white, or transparent.

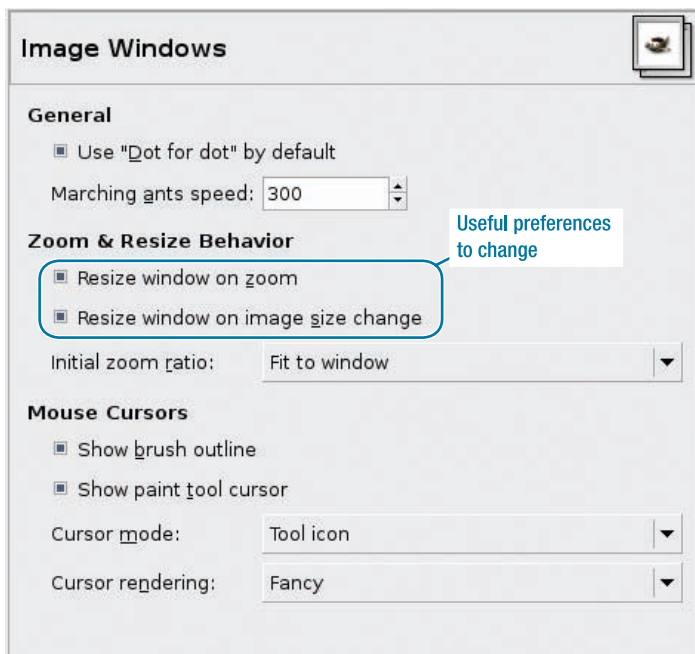
Finally, images can have a comment associated with them. In some image formats, such as GIF, the comment will be saved along with the image. You can use this to attach copyright notices or signatures to your images. Or you can keep the default, “Created with The GIMP” to let people know which image-editing program you like the best.

## Default Grid

Moving on past the *New Image* category, we find *Default Grid*. You can make the GIMP's grid visible from the *View* menu of an image window: *View* ► *Show Grid*. Grids can be helpful if you need to draw a lot of horizontal and vertical lines with well-defined intersections.

## Image Windows

The GIMP offers quite a few preferences that you can use to control the appearance and behavior of the GIMP's image windows (Figure 1-17). Use “*Dot for dot*” by default controls how windows are zoomed; you probably won't want to change it unless you regularly work with very high-resolution images intended for a printer. *Marching ants speed* controls the way selections are shown; you'll meet the marching ants a bit later, in Chapter 5.



**Figure 1-17.** *Image Windows preferences*

*Zoom & Resize Behavior* is a particularly useful group. Enabling these two preferences makes windows resize smaller or larger when you zoom in or out on the image. With the preferences disabled, the window will stay the same size when you zoom in or out, and will present scrollbars if needed. *Initial zoom ratio*, if set to 1:1, will always present images at full size when you first open them; otherwise, the GIMP will scale the image up or down to make it a comfortable size to fit on the screen.

*Mouse Cursors* lets you change the appearance of the cursor as you move your mouse over different parts of an image. In particular, the cursor can show the size and shape of the current

brush, and can reflect the current tool in use. Most people will want to leave these at their default settings, except perhaps on a very slow machine. However, some people may prefer the crosshair since it doesn't get in the way of the image as much as the regular cursor.

## Image Windows: Appearance

Remember all the parts of the image window, back in Figure 1-4? This preferences panel lets you configure which of those elements are shown. The “padding” mentioned on this screen is the blank space shown to either side of the image in an image window if you resize the window to be larger than needed for the image. You can also add or remove many of these items using the *View* menu inside the image window.

## Image Windows: Title and Status

This panel controls what information is shown in the title bar and status bar areas of image windows. You can use regular text, choose from a selection of useful samples, or design your own title or status string by selecting from special variables providing information about the current image.

---

**Note** You can change these settings interactively and watch what they do in the image window. For a list of these variables and how to use them, see the *gimprc* manual, discussed in Chapter 12.

---

## Display

When the GIMP displays an image that is partly transparent, the transparent parts are indicated with a light gray and dark gray checkerboard, so it's easy to see where they are. It may be easiest to think of this as GIMP keeping a gray checkerboard behind the image you're working on.

If the checkerboard annoys you, or if you're working on an image involving a gray checkerboard and can't tell it apart from GIMP's transparency indicator, you can change the way GIMP shows transparent areas here.

Also included in the *Display* preferences is the GIMP's notion of your display's resolution (in dpi). Normally, the GIMP can get this information from your windowing system, but if you need to change it, check *Manually* and enter the resolution here. Notice that there's another chain link icon here linking the horizontal and vertical resolution values together, as discussed earlier in the *Advanced Options* section of the *Default Image* category. You can also choose *Calibrate*, which brings up a window with rulers, and which you can measure and compare to a real ruler to get your monitor's true resolution. You may be surprised at how far off your computer's idea of your monitor resolution is. Be careful not to scratch your screen with the ruler! (A cloth measuring tape or a piece of paper can work well here too.)

## Color Management

These preferences let you specify color profiles matching particular monitors and printers. Most people should not need to change the color management preferences, but see Chapter 12 for more details on color management.



## Input Devices

The *Input Devices* preferences control goodies such as graphics tablets, which you can use in addition to a regular mouse. If you have such a tablet, choose *Configure Extended Input Devices* to select which of your tablet's tools will be used by the GIMP.

---

**Tip** The *Save Settings* button in this preferences pane is useful even if you don't have a graphics tablet. If you have a tool you use frequently in the GIMP, you can make the GIMP select that tool automatically whenever you start up by selecting the tool in the Toolbox and then clicking *Save Input Device Settings Now*.

---

## Input Controllers

Here, you can configure the GIMP's behavior in response to hardware such as a mouse wheel when used in a GIMP image window.

By default, the mouse wheel by itself scrolls up and down in the image window (assuming a scrollbar is present); holding the Shift key down while using the mouse wheel scrolls left and right. Holding the Control key down makes the mouse wheel zoom the image larger and smaller. Other combinations of key modifiers cause other behaviors: in this preferences pane, you can view the current settings, delete current settings, or add your own by clicking on a setting, such as *Scroll Up* (Shift+Control) then clicking the *Edit* button. This will bring up a dialog showing the actions you can trigger using the mouse wheel.

Clicking on the *Main Keyboard* tab shows functions that are bound to the arrow keys on the keyboard. By default, the arrow keys scroll by a small amount, unless a tool is selected (such as the Move tool) in which the arrow keys do something else. Again, you can customize the behaviors in response to the arrow keys using the *Edit* button, or *Delete* a behavior if it bothers you.

## Window Management

This group of preferences controls the behavior of some of the GIMP's windows. If you have problems such as windows not staying on top when you think they should, or windows starting up in the wrong position, try changing these preferences.

*Window Manager Hints* controls whether the Toolbox and other docked dialogs such as the Layers dialog should be treated like dialogs and always kept above image windows, or treated like normal windows. Choosing *Utility Window* keeps the window always on top of every GIMP image window; *Keep Above* tries to keep it above every window on the system. Not every system may support all of these behaviors.

*Activate the focused image* makes an image window active as soon as it receives focus: in other words, as soon as you click on it, tab to it, or sometimes just move the mouse over it. Moving to a new active window changes the behavior of other open dialogs, such as the Layers dialog, which can be inconvenient. Disable this preference if you have this problem.

If you do disable *Activate the focused image*, you may find that you need a way to make a new image active without making any changes to it. An easy way to activate an image is to focus the image window, then hit the space bar.

The remainder of the *Window Management* pane concerns the GIMP remembering its window positions.

## Folders

The *Folders* preference category lets you configure where the GIMP stores its various preferences, plug-ins, brushes, and other information it needs. Normally, it's best to leave all of these at their default values. When you start adding collections of brushes and plug-ins, you can add folders to the lists here to make it easier to keep track of what's where.

## A First GIMP Project

Enough exposition—the only way to learn the GIMP is to edit images, so let's try making an image!

Begin with any photo you like. I'll start with a vacation photo of some ancient ruins from Chaco Canyon in New Mexico.

You could use an image you already have in a file on your hard disk, and use the GIMP's *File* ► *Open...* dialog to open it. But you may want to choose an image you're keeping on a website. The GIMP makes it especially easy to open an image that is on the web: just drag from the image in the browser window, and drop it onto GIMP's Toolbox window. Drop it anywhere over the tool buttons; it isn't important which button is underneath.

When you drop the image (release the mouse button) over the Toolbox window, a New Image window containing the chosen image opens (Figure 1-18). At this point, you don't yet have a copy of the image on disk; that won't happen until you save the image and give it a file name.




Figure 1-18. Image window showing Chaco Canyon ruins

To jazz up the image, it's time to introduce a new character: Wilber (Figure 1-19). Wilber is the GIMP's mascot, drawn by GIMP contributor Tuomas Kuosmanen ("tigert").



**Figure 1-19.** *Wilber, the GIMP mascot*

Images of Wilber can be found on the GIMP's website, <http://www.gimp.org>. Try dragging one of these images from your browser window into the image window you already opened. The image is added to the current window as a separate layer (Figure 1-20). If dragging doesn't work, save the image from your browser (right-clicking on the image will usually offer that option), and then drag it from your file manager into the GIMP image window.



**Figure 1-20.** *Chaco Canyon image with Wilber added*

---

**Tip** When you're inserting a character like Wilber into another image, you'll usually want the Wilber image to have a transparent background, so you don't end up with a white square around the image you inserted. You'll learn more about ways to make a colored background transparent in Chapter 5. For now, just choose an image in GIF or PNG format that already has a transparent background. Fortunately most of the Wilber images on the GIMP website will work fine.

---

If you have your Layers dialog open, you may notice the new layer has appeared. If not, don't worry about it; we'll talk about layers in Chapter 3. The yellow-and-black square around Wilber is the layer boundary: it shows the size of the layer. It's not really part of the image, and will not be visible to anyone else. Wilber appears in the middle of the image, which probably isn't where you want him. In my Chaco Canyon image, I want Wilber's head to appear in the window. To move a layer, select the Move tool from the Toolbox window (Figure 1-21).



**Figure 1-21.** *The Move tool*

With the Move tool selected, you can drag a layer around with your mouse to position it exactly where you want it. Then save the image by choosing *File* ► *Save As...* from the image window, and give it a file name, such as *chaco-wilber.jpg*. (You'll learn more about image file formats in Chapter 2.) If the GIMP pops up a dialog about JPG not handling transparency or layers, just click OK.

Figure 1-22 shows the final image.



**Figure 1-22.** *Wilber visits Chaco Canyon*

That's all there is to it! Now you can invite Wilber (or any other character you like) along on your next vacation.

## Summary

At this point, you should have a basic understanding of the GIMP's windows, menus, and preferences. You've seen how to combine two images in a simple GIMP project, and perhaps you've spent some time playing with some of the GIMP's built-in filters and plug-ins. Now let's move on to the details of image editing and explore how to make digital photographs look better.



# Improving Digital Photos

**Y**ou have your shiny new digital camera, and it takes great pictures. But straight out of the camera, the photos are huge, and far too big to put on a website or email to friends. Maybe they have other problems, too, which you'd like to correct before you show them to anyone.

This chapter will explore some of the most common ways you can use the GIMP to improve your digital photos and share them with the world. It will cover the following:

- Opening files
- Rescaling
- Saving files
- Cropping
- Brightening and darkening
- Rotating
- Sharpening
- Fixing red-eye

## Opening Files

The first step in editing any image is to open it.

In Chapter 1, you saw that you can drag an image from a web browser or file manager to the GIMP's Toolbox window to open it. In addition, you can drag one or more images to the GIMP's desktop icon, if you have one, or simply run GIMP from the command line: `gimp file1.jpg file2.jpg . . .`

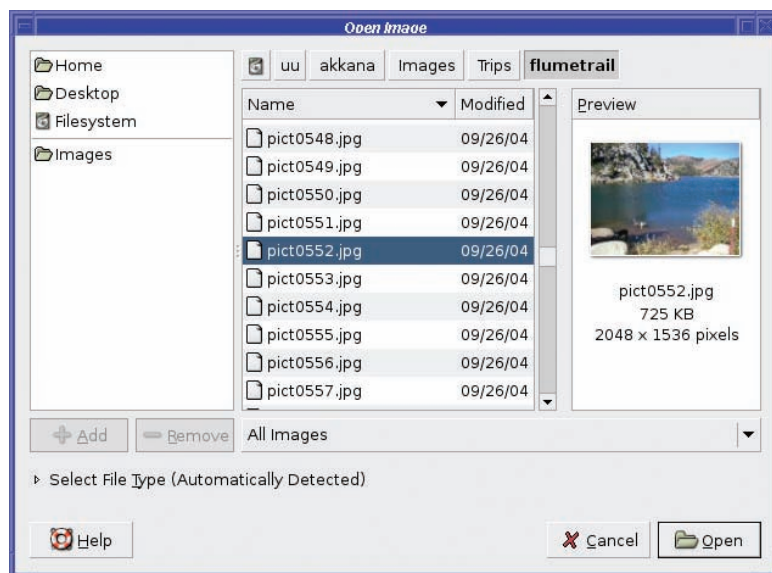
---

**Caution** Some of these suggestions, including drag and drop, may not work properly on the Macintosh version of the GIMP. On a Mac, your best bet is to drag files onto the GIMP icon in the Dock.

---



Once the GIMP is running, you can open additional images with the *File* ► *Open...* dialog, which is shown in Figure 2-1.



**Figure 2-1.** The GIMP's *Open* dialog

The GIMP's *Open* dialog lists files and folders in the main frame of the dialog. Above the file area is a sequence of buttons representing the filesystem path to the currently chosen folder. Clicking on these buttons can move you up one level, or several.

You can select more than one file at a time in the file list. Click on the first file, and then hold down the Shift key while clicking on another file to highlight a range of files; or hold the Control key down while clicking individual files to select them. The GIMP will open each file as a separate image.

Try typing the first few letters of an image's name to jump to that image in the file list. Control+L will bring up a dialog that allows typing or pasting a file name. On Linux, typing certain characters (such as /) will cause this dialog to pop up without the need for Control+L.

On the left side of the dialog is a list of *bookmarks*. You can double-click on a bookmark to jump directly there, or add the current folder to the bookmarks list by clicking the *Add* button.

Underneath the file list is a menu for file type, initialized to *All Images*. You can use this menu to restrict the types of images you see in the dialog.

Below the *Add* and *Remove* buttons is an expander for *Select File Type*. This is for unusual situations where the GIMP cannot detect the format of an image file automatically. It is seldom useful: if the GIMP can't detect an image file's type on its own, it usually means that the file has been damaged in some way, in which case this selector won't help. However, if you have an image that you think may have the wrong extension, it's worth a try.

Besides the *Open* dialog, GIMP offers two more *Open* options in its *File* menu: *Open Location*, which allows you to paste a web location, and the *Open Recent* menu, which lists images you've edited . . . well, recently. If the file you're looking for is no longer listed in the *Open Recent* menu, you may still find it in *Document History* at the bottom of the *Open Recent* menu.

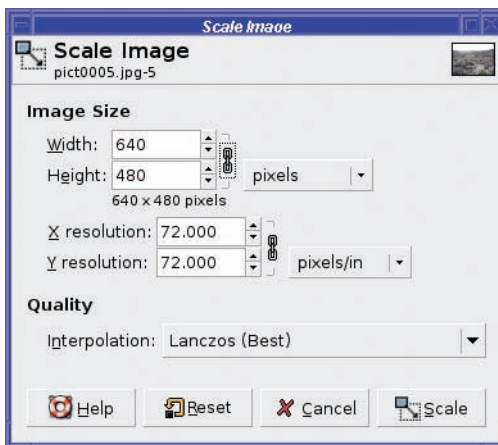
**Caution** When editing a digital camera photo, it's a good idea to make a copy of the image first, and then edit the copy. Although you can use the GIMP's Save As dialog to save a file to a different name, it's easy to forget to do that, and overwrite your original image.

## Rescaling

Modern digital cameras take beautiful photos in high resolution: 3, 4, 5 megapixels or even more. That's great, but higher resolution makes for larger files of at least half a megabyte each and perhaps much more. If you put images this big on the web, or send them via email, the recipients will not thank you.

The solution is to *scale* the images: make them smaller. Don't confuse this with *resizing*: in the GIMP, *resize* means to change the size of the image without changing the contents. Resizing an image larger just creates more blank space around it (enlarging the *frame*), while resizing it smaller eliminates part of the image. To take the image contents and blow them up or shrink them down, scale is what you want.

The Scale Image dialog (Figure 2-2) is accessed from the *Image* menu in the image window: *Image* ► *Scale Image*.



**Figure 2-2.** The Scale Image dialog

The dialog has three sections. The most important section comes first: the new size (width and height) of the image. By default, size is given in pixels. This is a good choice, and you should use it unless you have a firm reason to do otherwise.

The width and height are “chained” together, as discussed in the section “Default Image” in Chapter 1. You can type a new number into either the *Width* or *Height* field, then hit the Tab key or click in another field, and the other dimension will automatically update to keep the *aspect ratio*—the ratio of width to height—of the image unchanged. This prevents the contents of your photo from getting taller and skinnier, or shorter and fatter. Very convenient!



If you actually do want to scale the two dimensions independently, clicking the chain link icon breaks the chain and unlinks width and height, so you can change one without changing the other. Welcome to the Hall Of Funhouse Mirrors. Clicking on the chain link again will link them back together.

Even if you work in pixel dimensions, you may want to use the units drop-down menu (initially in pixels) in order to scale an image to some multiple or fraction of its current size. With the units menu set to *percent*, you can set the width or height to 2 to make the image twice as large, or to .5 to make it half the size. (You only need to set width or height, not both; if they're chained together, the other will adjust automatically.)

### THE UNDO COMMAND: A GIMP USER'S BEST FRIEND

If you're curious about how the chain link icons work, or how the image might look when scaled to a particular size, you can try it without any fear of losing your data. The GIMP's Undo command is your friend! Get familiar with it early on. It's very useful to be able to try things while knowing you can always undo your changes and go back to the previous image.

Undo is in the *Edit* menu, along with *Redo* and an item called *Undo History*, which brings up a window showing you the history of commands you've used recently. You can also undo the last command by pressing Control+Z. Undo is so frequently useful that it's well worth learning that keyboard command.

The GIMP can undo several operations in a row, not just the last single operation. By default, it remembers the last five operations. You can change that number in Preferences (*File* ► *Preferences*, under *Environment*). A larger Undo stack takes more memory, so it may be best to stick with the default setting unless your machine has quite a lot of RAM. There's also a preference to limit the amount of memory the Undo system will use, though the preference for the number of Undo operations takes precedence, so you know you'll always have at least the number of Undo operations you specify in Preferences.

Try Undo now: type something into the Scale dialog, perhaps even with width and height unchained from each other, and click the Scale button. Then undo the Scale operation, using Control+Z or Edit ► Undo, so that you're confident it works.

## What Size Should You Choose for Your Images?

For photos intended for web pages, it's usually best to keep the largest dimension of the image at 640 pixels or less (400×300 images are free on eBay!). If you have a lot of images, or if many of the people who will view your images have slow network connections, smaller is better.

The *Resolution* section in the Scale Image dialog allows you to specify the number of pixels per inch (or other units, if you choose). This doesn't change the size of the image; it's merely a note that gets added to the image so the image size can later be converted to physical units (such as inches). Most of the time, you won't need to look at these values.

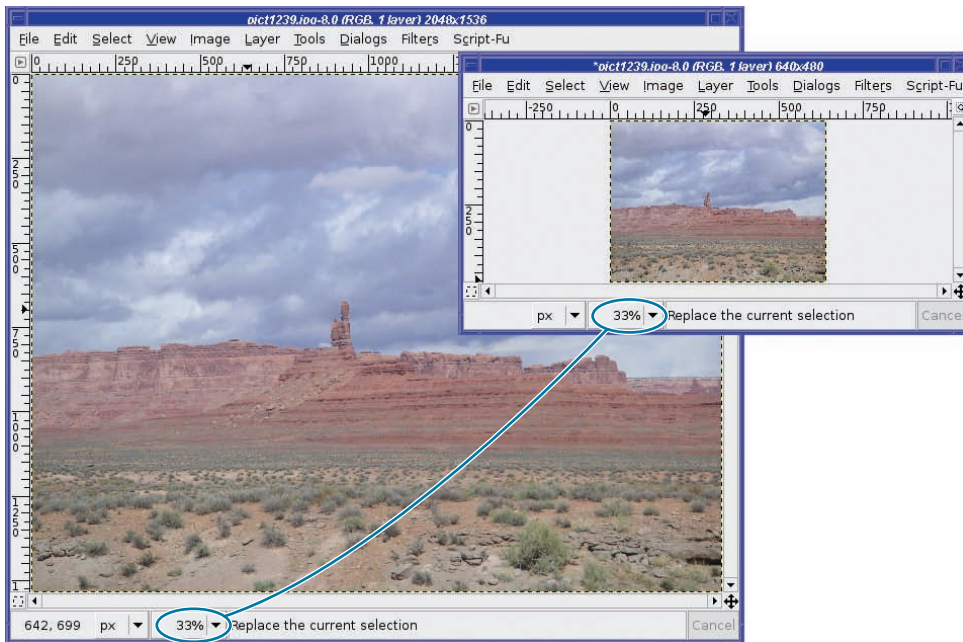
The *Interpolation Quality* controls details of how the image is scaled. *Linear*, which is the default in at least some versions of the GIMP, does not produce the best quality, but it's fairly fast. Unless you're on a very slow machine, you'll probably want to change this to one of the higher-quality settings, such as *Cubic* or *Lanczos (Best)*. These settings will run a little slower, but they'll produce much better-looking images, especially if you ever scale images larger instead of smaller. I recommend changing this value permanently, in the Preferences dialog under *Tool Options*, so you don't have to think about it when you're actually using the Scale dialog.

---

**Note** Which quality setting is considered “Best” can vary with the GIMP version. However, you’re always safe choosing whichever option is labeled “(Best).”

---

Usually, when you click *Scale* in the dialog to make a digital photo smaller, and the GIMP rescales the image, the window will resize to be tiny (Figure 2-3). Don’t be fooled by this: the image probably isn’t as small as the GIMP is showing you.



**Figure 2-3.** *The image window is now too small!*

What happened? Since the original image was too large to fit on the screen, the GIMP showed it as “zoomed out”—in the example in Figure 2-3, it is shown at only 33% of actual size. After the image is scaled, the GIMP keeps the same zoom setting. So now, even though the image could easily fit on the screen, you’re still seeing it at 33%. The solution is to zoom after rescaling. Use the *Zoom* menu in the image window, the menu under *View* ► *Zoom*, or the keyboard shortcut **+**, or type **1**. Typing **1** (the numeral one) in the image window always zooms it to full size (100%). Typing **+** (plus) zooms in (makes the image appear larger); **-** (minus or dash) zooms out. When preparing images for the web, you’ll probably want the image window at 100% so you can see exactly what your audience will see.

---

**Caution** The + and – keys on the numeric keypad may not work. If they don't, use the ones on the regular keyboard. On some earlier versions of the GIMP, zooming in used = instead of + (on most US keyboards, = is the unshifted version of +, and so is easier to type). Remember, you can always change these key bindings, as described in Chapter 1 in the section “User Interface and Keyboard Shortcuts.”

---

Scaling an image is easy, and it's an important operation when you're sharing images with other people. For many images, scaling may be the only operation you need. But if you plan other operations that may affect image size, such as rotation or cropping, do them first. Scaling should usually be the *last* thing you do when you're editing an image; that way, you have the greatest control over the image's exact final size.

The exception that proves the rule is a very slow machine. If you scale the image first, subsequent operations will go much faster. However, you'll find that this matters only with *very* slow machines—or *very* large images!

## The Scale Tool

Most of the time, *Image* ► *Scale* is perfect for rescaling an image. But there's another way of scaling: the Scale tool.

You'll find it in the Toolbox next to the Rotate tool, or through the menu *Tools* ► *Transform Tools* ► *Scale*. This tool lets you scale interactively: with the tool selected, click anywhere in the image to activate it, and then drag from any corner of the image into the image.

The Scale tool is a bit trickier to use than *Image* ► *Scale*... for a couple of reasons. First, you have to remember to click in the tool to start the operation. Second, it only scales the current layer, not the whole image, and it doesn't change the image size at the end.

Third, by default it doesn't maintain the layer's original aspect ratio. You may end up with a “funhouse mirror” version of your image where everything looks tall and skinny or short and fat. *Keep aspect* in the tool options can fix that (there are also options to keep height while varying width, and vice versa).

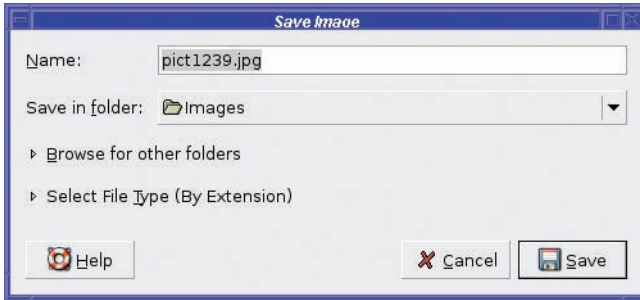
For digital photo operations, you probably won't want to use the Scale tool. But keep it in mind for later when you're using lots of layers and doing crazy things to them. It might be just the ticket then, as it does have the advantage of giving a live preview of what the image will look like when scaled. Sometimes that's awfully helpful.

## Saving Files

Image editing is just like any other type of document editing: once you've made any significant change to a file, you should save it. It's smart to save your work frequently, in case of a computer crash or other disaster.

The GIMP uses Control+S as the keyboard shortcut for *File* ► *Save*, just like most word-processing programs. If the image already has a file name, that's a quick way to save your work.

If you've edited an original image, modified it, and now want to save to a new file name, use *File* ► *Save As...* (Figure 2-4). You'll also see this dialog the first time you save a new image.



**Figure 2-4.** *The Save dialog*

The *Save in folder* popup offers a list that corresponds to the “bookmarks” you’ve set up in the GIMP’s Open dialog. You can get a file selector that gives you more choices by clicking on the expander next to *Browse for other folders*.

The *Select File Type* expander lets you choose an image format. Normally, you won’t use it: just type the extension as part of the file name (e.g., *mypicture.jpg* to save in JPEG format) and GIMP will use the appropriate format. The expander is there in case you need to see a list of formats the GIMP can handle, or in the rare case that you need to save a file with an extension that doesn’t match the file type you need to use—for example, if you want to save a JPEG image but name it “*image.gif*” instead of “*image.jpg*”. You’ll learn more about file types in the next section, “Image File Types.”

Once you *Save As...*, the GIMP considers your entry to be the new name for the image. From then on, you can save and any changes will be saved to the new file name, without overwriting the original file. The file name is visible in the title bar of the image window, so you can check easily to make sure you’re saving to the right file name. The GIMP has one other save option worth noting: *Save a Copy...* This works like *Save As...* in that you can choose a file name different from the one you’re currently using. The difference is that *Save a Copy...* does not change GIMP’s notion of the current image file name. After you *Save a Copy...*, subsequent saves will still use the original file name. This is very useful when you’re making repeated edits to an image, and you want to save it in a format that preserves all your work, such as XCF; but you want to share the final result on the web in a more efficient format, such as JPEG or GIF.

## Image File Types

The choice of file types used to hold images—GIF, JPEG, BMP, PNG, TIFF—is baffling. What do each of these file types mean, and which format should you use?

When choosing a format for the web, or to send images by email, the most important criterion is usually file size. In order to view the image, each user will need to download it. This is especially important if anyone in your intended audience might be connected via a slow modem, or if you’re sending images by email and some recipients might have an email account with a disk quota. Also, many email systems will simply refuse to transfer files that are too big.

The other important criterion when choosing an image format is image quality. There’s typically a tradeoff between file size and quality: you can compress an image to a smaller size only by sacrificing some quality.

## JPEG

For sharing full-color photographs, the best format to use is usually JPEG. Pronounced “jay-peg,” the name comes from the group that defined the format, the *Joint Photographic Experts Group*. Usually, a JPEG file will have an extension of “.jpg”. JPEG images are highly compressed and are encoded in full color, so it’s a very efficient format for photographs or other images with a wide range of colors.

The drawback to this format is that the compression is *lossy*. That means that every time you read a JPEG file, make a change (however minor), and write the file back to disk, the image quality degrades slightly. Don’t use JPEG for images you plan to edit over and over. But as a format for exchanging photographs with other people, JPEG excels. (There’s no quality loss when you copy the file from one place to another, only when you edit it and write it back to disk.)

You may be wondering, “My camera stores images as JPEGs. Does that mean I’m losing quality?”

The answer is yes. That’s why many cameras provide a non-lossy *raw* format (discussed in the section, “Other Formats”) as well as JPEG. However, JPEG files are so much smaller that you can store many, many more images than is possible with raw mode on the same memory card. Use raw format if you’re worried about it, or if you’re bothered by the image quality of a JPEG; but for most people, the slight quality loss caused by using JPEG on a camera, plus editing an image once or twice and saving it *again* in JPEG format, will not be noticed. The space savings are worth it.

## GIF

The *Graphics Interchange Format*, pronounced “giff” with a hard “g” as in “graphics,” is an *indexed* format. This means that it uses a fixed list of colors instead of encoding every color separately. This is very efficient for images with a small number of colors, like a five-color corporate logo.

GIF can represent up to 256 colors (256 is 2 to the 8th power, so this is also called *8-bit color*). Typical photographs have many more colors than that, so saving a photograph in GIF format will usually result in very poor quality. Even worse, with 256 colors, the file will usually be larger than a full-color JPEG version of the same image! The lesson is clear: don’t use GIF for photographs, only for simple icons and logos.

The GIF format offers two very useful features: transparency and animation. With transparency, you can make an icon with a clear background. If you display it on a web page, or on a button in a program’s user interface, whatever’s behind the icon (even if it’s another image) will show through. GIF doesn’t allow for partial transparency; a pixel is either fully transparent or not transparent at all.

GIF animation allows you to create images that move. Most web browsers support animated GIF images. (Whether this is a good idea is unclear; some users dislike animated images on web pages, and will avoid websites that use them. But they have their uses.) To learn how to create GIF animation, see “Making Simple GIF Animations” in Chapter 3.

## PNG

PNG, pronounced “ping,” once stood for “PNG’s Not GIF,” though it’s now sometimes referred to as the Portable Network Graphics format. It’s a relatively new format, originally intended as a replacement for GIF because of legal issues regarding the GIF format.

PNG can be used for full-color images, like JPEG, or for indexed images, like GIF. In full-color mode, it’s not nearly as efficient as JPEG; a PNG will be much larger than a JPEG of the same image. However, it’s not lossy: if you read a full-color PNG, make a modification, and write it back out, the quality will be just as good as it was before. This makes full-color PNG a very good format for storing original copies of your own images if you might want to edit them later.

Indexed PNG is just as efficient as GIF (sometimes more efficient) for images such as logos and icons with a small number of colors, and it can support more than 256 colors. Nearly all web browsers support PNG, so it’s safe to use PNG images on web pages. PNG supports transparency, like GIF, and it also supports partial transparency so you can have translucent areas. However, some browsers (such as Internet Explorer) don’t support PNG transparency, so transparent PNG images on a web page may not always display properly.

PNG doesn’t support animation. There’s a format called MNG that adds animation to PNG images, but no web browsers support it yet. For animated images, GIF is pretty much the only choice.

## XCF

This is the GIMP’s own format. When you’re editing an image and you have a lot of layers and paths and other information you want to save, this is the format to use. Otherwise, use some other format. XCF files are usually quite large, and can only be read by the GIMP, not by other programs. Since XCF files are so large, you can compress them using the GZIP or BZIP2 formats—filename.xcf.gz or filename.xcf.bz2—and GIMP will handle the compression and decompression when it reads or writes the file.

## TIFF

*Tagged Image File Format* (with a file extension of either .tif or .tiff) is another full-color, non-lossy format. Like full-color PNG, it’s not very compact (don’t use it on web pages—many web browsers can’t display it anyway), but it’s fine for keeping originals of images you might want to edit again. The reason I recommend using PNG instead of TIFF, aside from web-browser compatibility, is that TIFF isn’t a single standard, but many different standards with different interpretations. A TIFF written by one program may not always read correctly in another program. The GIMP will read most TIFF files, but sometimes it will “have a tiff” about certain types of TIFFs.

One advantage to TIFF over other full-color formats is that it handles a wider range of colors. For the technically-minded, it can handle 16 bits per color channel. This is important to some professional artists and graphic designers. However, the GIMP doesn’t handle 16-bit color anyway (that’s planned for a future release), so if you’re editing in GIMP, you don’t get any benefit from using TIFF.

## Other Formats

There are more image formats. Lots more. *Too many* more. Here are some of the most common:

- *Raw*: Not actually a format, *raw* is a term encompassing all the various proprietary formats used by camera manufacturers. The GIMP supports some of these formats (usually requiring external plug-ins); other formats are held as “trade secrets” by the camera manufacturer, and your only option is to use another program to convert the image to a more standard format before editing it.
- *BMP*: This is Microsoft’s Windows Bitmap format. BMP files are quite large and don’t offer any advantage over PNG or TIFF. Save to BMP if you need to (for example, for a Windows icon), but otherwise it’s usually better to choose a different format.
- *PCD*: This is Kodak’s proprietary Photo CD format. It includes several resolutions within one file (so files tend to be very large) and is not lossy. The GIMP doesn’t handle PCD directly, though there’s a plug-in available to read it. You’re usually better off converting PCD files to something else.
- *PSD*: This is Adobe’s proprietary Photoshop format. It saves layers and other information, analogous to XCF in the GIMP. PSD is really two formats: the GIMP can read the older version, while the newer one is a closed standard that can’t be read by the GIMP or most other programs.
- *ICO*: This is the Microsoft Windows Icon format. It can contain several resolutions in one file. This format is useful not only for Windows icons, but also to create a *favicon* for your website. These show up if someone bookmarks your site, as well as next to the web page address in some browsers.
- *PDF and PostScript*: These are vector graphics formats, not raster (pixel) graphics like the other formats discussed so far. Instead of representing an image as a rectangular collection of pixels, a vector image is a collection of drawing instructions involving points, lines, and curves. The GIMP can’t edit vector graphics directly, but it can import PostScript or PDF through a plug-in that converts the image into a raster image. If you’re planning to save your image as PostScript or PDF, you’re usually better off using a program intended for editing vector graphics.
- *SVG*: Scalable Vector Graphics is another format that is growing in popularity on the web. The GIMP can import images from SVG, and in the future it may be able to export them to SVG as well. In addition, you can use the Path tool (which you’ll meet in Chapter 5) to make simple vector graphics that you can save as SVG, and you can import paths from SVG.

## Experimenting with JPEG and GIF Settings

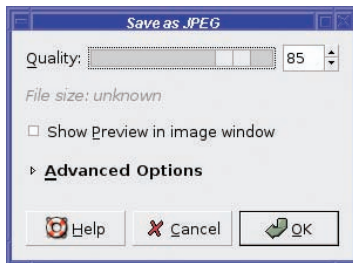
Some image formats let you adjust settings to trade quality for smaller file size. In some cases (for instance, small buttons on a web page) it may be important to make the file as small as possible, and a slight loss of quality may be worthwhile.

JPEG and GIF both have quality adjustments, but their methods are very different.




## JPEG Quality Settings

JPEG's quality adjustments come by compressing the image more or less thoroughly. JPEG's compression takes advantage of tricks of human perception; it tries to degrade quality in ways that won't be very noticeable. By experimenting with the *Quality* setting in the Save as JPEG dialog (Figure 2-5), which appears as soon as you've specified a file name ending in .jpg, you can find a balance that gives you a very small file size but still a pretty good-looking image.



**Figure 2-5.** JPEG save options

Begin by making sure the checkbox for *Show Preview in image window* is checked. This lets you see the effect of changing the quality setting. The file size is also shown in the dialog. The GIMP's default quality setting is 85, which is a good choice for photographic images stored locally; but you may find that you can get away with settings as low as 50 with surprisingly little degradation. Figure 2-6 shows an example of how the quality can change with different compression settings, but I encourage you to try this on one of your own images, to see the differences in quality and file size.



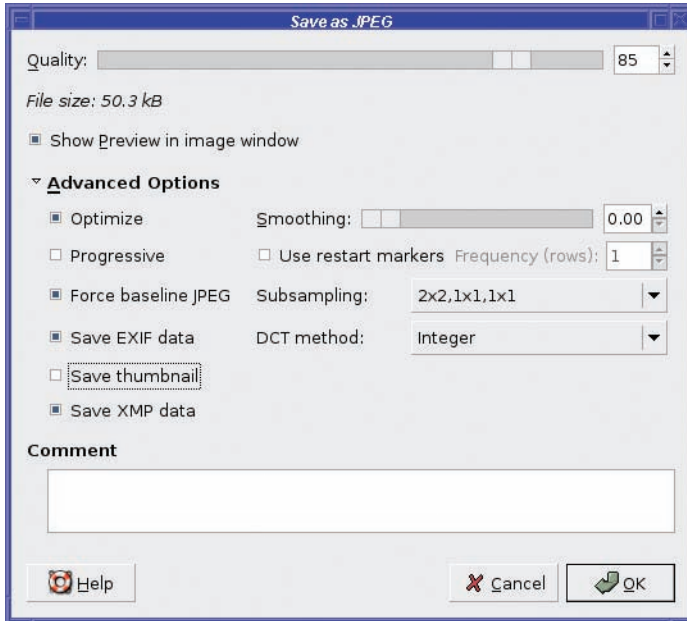
**Figure 2-6.** Comparing JPEG quality settings

You may be tempted to move the slider to 100% for images you store locally. Surely that would be best for images you're storing, and would prevent loss of data from JPEG compression?



Alas, even at 100%, JPEG still loses some data every time you save; using a setting of 100% will produce a file two or three times larger than using a setting of 95%, and the quality won't be any better. There's really no advantage in using quality settings above 95%. If you want to save without any loss, you're better off using formats such as PNG or TIFF.

Finally, JPEG has some other settings in the *Advanced* area of the Save as JPEG dialog (Figure 2-7).



**Figure 2-7.** Advanced JPEG settings

*Optimize* gives you an additional reduction in file size without any further reduction in image quality. It's on by default, and there's no reason to change that.

*Progressive* is a useful setting for images that will be uploaded to the web. It makes the image load in a different way, so people viewing the image will see a poor-quality version right away, which gradually improves, instead of seeing the image load line by line starting from the top.

*Smoothing*, *Use restart markers*, *Force baseline JPEG*, *Subsampling*, and *DCT method* control details of the JPEG format. You shouldn't need to change these unless you want very fine control over the way the JPEG file is saved.

*Save EXIF data* is a good option to know. The JPEG format includes a way of holding information about the file in a format called *Exchangeable Image File Format*, or EXIF. Most digital cameras add EXIF information about the date the photo was taken, the resolution, and the camera's settings such as lens focal length and whether or not flash was used. The GIMP cannot show or edit this EXIF information (that will probably be available in some future release), but some versions can preserve the information so that you can view it with other programs. If you don't want the EXIF information preserved, deselecting this box will make the file slightly smaller. (You can see exactly how much smaller by watching the *File size* value at the top of the dialog.)

In GIMP 2.4, this has been replaced by *Save XMP*, which can save some but not all of the EXIF information. Full EXIF handling may be restored to the GIMP in a future version.

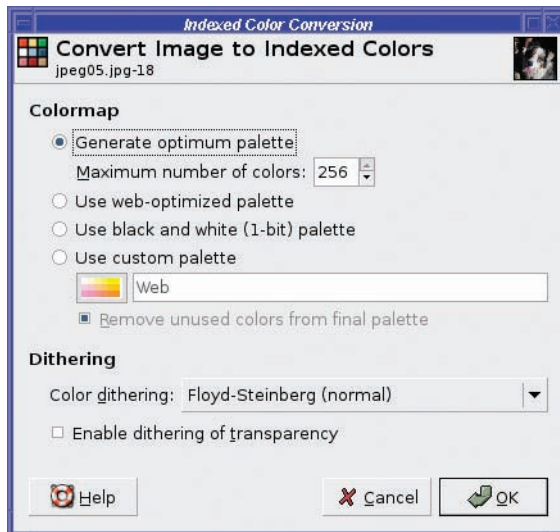
EXIF can embed a thumbnail of the image in the file. Many cameras include this thumbnail, which adds several kilobytes, but most programs don't do anything useful with the EXIF thumbnail. Unchecking *Save thumbnail* prevents saving this extra information, making the resulting file smaller.

*Comment* is a place to put text you might want to add to an image. You can use this to include your name, a copyright notice, or details about where the photo was taken or what it portrays.

## GIF and Indexed PNG Quality Settings

GIF and indexed PNG quality settings are quite different from JPEG. Since these indexed formats use a fixed number of colors, you can make the files much smaller by reducing the number of colors they use.

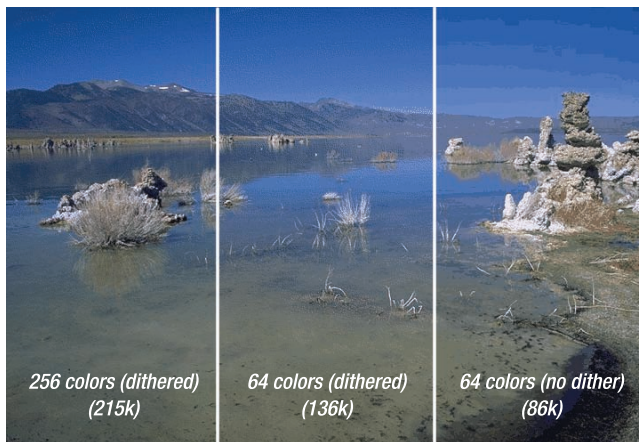
To do this, you need to know about *image modes*. Most images edited in GIMP are full-color images, denoted as *RGB* in the title bar of the image window (back in Figure 1-4). But images that are destined to be saved as GIF or indexed PNG should be converted using the *Image* ► *Mode* ► *Indexed...* menu item, which brings up the dialog shown in Figure 2-8.



**Figure 2-8.** *Convert Image to Indexed Colors dialog*

When converting an RGB image to indexed mode, you first need to choose a *palette*, the set of colors that will be represented in the final image. You can choose a predefined palette, such as a *web-optimized palette*, or one chosen from the custom palette menu; you can choose a black-and-white palette, if you know that the image contains only those two colors; or, the most common case, you can let the GIMP create a palette based on the colors actually in the image. You specify the number of colors, up to a maximum of 256, and the GIMP will try to approximate any other colors that exist in the image by using combinations of the colors in the palette.

This process of approximation is called *dithering*: it involves combining pixels of several different colors. You can choose from several different types of dithering that create different effects, or you can specify no dithering, which may create a smaller file size and a cleaner appearance for images with sharp lines and only a few colors. If the image contains translucent areas, the GIMP can also use dithering to simulate this in formats such as GIF, which don't support partial transparency (Figure 2-9).



**Figure 2-9.** Different levels of indexing, showing file sizes when saved as GIF. Note that all the GIF versions of this photo, even the smallest, result in a bigger file than the best JPEG.

The Convert Image to Indexed Colors dialog does not include a *Preview* button. So you'll probably need to run it several times: choose some settings, click OK, and then use Undo to go back to RGB mode and try again with different settings. It's worth spending some time on this if you're trying to squeeze an indexed GIF or PNG image as small as possible.

---

**Note** The GIMP offers a third image mode, *grayscale*, for black-and-white images. Grayscale mode is useful for scans of black-and-white documents or photographs, or simply for converting color photos to black and white.

---

When saving an RGB image as GIF, if you skip the step of converting it to indexed mode, the GIMP will offer to do the indexing for you. In this case, the GIMP will choose a palette intended to represent as many of the image's colors as possible. This helps preserve image quality, but it doesn't do much to reduce the file size. You may even end up with a larger, yet poorer quality file than if you'd used JPEG. You're much better off doing your own conversion.

The Save as GIF dialog offers a few additional options.

- *Interlace* is like JPEG's progressive option: it makes images load in a different way, so that someone viewing them in a web browser can see more of the image before it has loaded completely.
- *GIF comment* is like a JPEG comment, a place where you can store copyright information or details about where or how the image was taken. (GIF doesn't include EXIF, so the GIF comment is the only place such information can be recorded.)
- *Animated GIF Options* controls ways of making an animation, as you'll see in Chapter 3.

Saving an RGB image as PNG will save it as a full-color PNG. If you want indexed PNG, you *must* convert the image to indexed mode before saving it. That said, the Save as PNG dialog offers yet another set of options.


- *Interlacing* is similar to GIF's interlacing or JPEG's progressive options.
- *Save background color* relates to images with transparency. Since not all web browsers handle PNG transparency, those that don't can use the GIMP's current background color instead.
- *Save gamma* saves information about your monitor. Some viewing programs can compensate for this, and try to adjust the PNG image so that it looks the same on the viewer's monitor as it did on yours. (Most viewing programs probably won't get this right, so for most people, there's no point in checking this option. It increases file size very slightly.)
- *Save layer offset* is only meaningful if you're saving a layer that has been moved relative to the rest of the image (you'll learn about layers in Chapter 3). Most viewers probably won't handle this correctly either.
- *Save resolution* records the current resolution (dots per inch or dpi) of the image, and is only useful if you're trading files with someone whose software cares about dpi.
- *Save creation time* records the time the image was last modified.
- *Save comment* is like the GIF or JPEG comment, but alas, you can't set it in this dialog.
- *Save color values from transparent pixels* is similar to *Save background color*, intended mainly for browsers that don't handle transparency properly.

Finally, PNG offers a settable *Compression level*. Unlike JPEG, PNG's compression is not lossy. The only disadvantage to using high compression levels is that it will take slightly longer to save the file. Most of the time difference happens when saving; it doesn't take a viewer much longer to uncompress a highly compressed PNG image. So usually it's best to use the largest amount of compression possible.

## Cropping

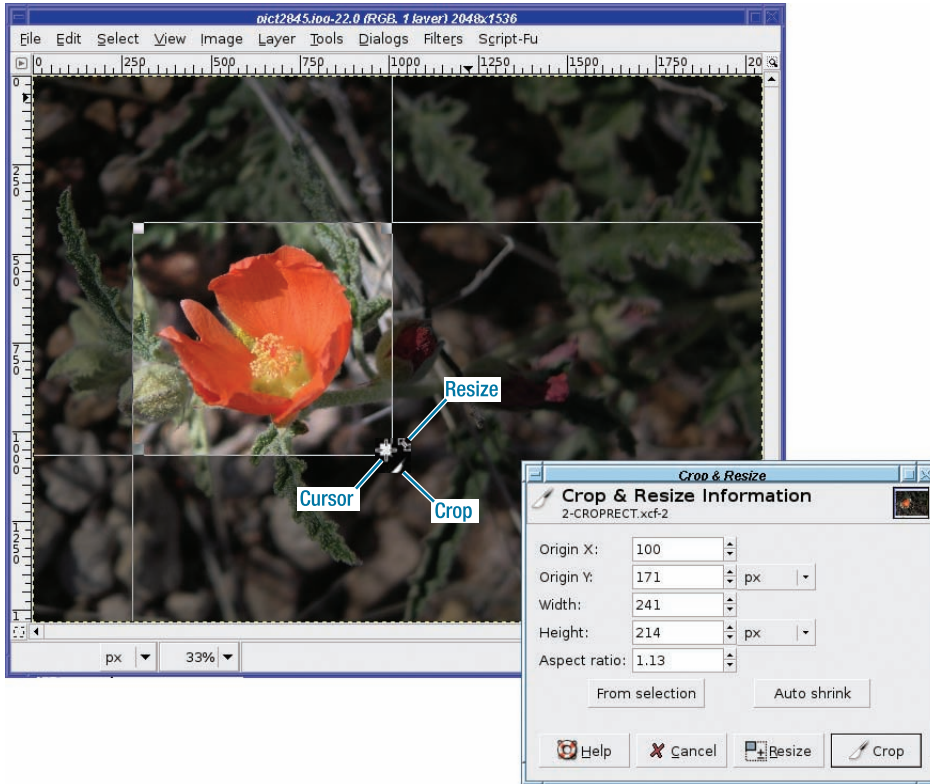
Fairly often, a photo has more in it than you want. Perhaps you took a picture of something far away, and your zoom wasn't quite enough to get "up close," so the subject only occupies a small part of the center of the photo. Or perhaps there's some unwanted object intruding on part of the frame, such as another person, a garbage bin, or your finger. (Not that you'd ever take a photo with your finger in front of the lens! At least, with the GIMP, you don't ever have to admit to doing such a thing.)

The solution to this problem is *cropping*. Click on the *Crop tool* button in the Toolbox to select it; it looks like a scalpel blade (Figure 2-10).



**Figure 2-10.** *The Crop button in the Toolbox*

To use the Crop tool, press the left mouse button down in the image window at the upper-left point of where you want the cropped image, drag down and right to where you want the lower-right point to be, and then release the mouse button. (You can actually start at any corner, as long as you drag diagonally to the opposite corner of the rectangle you want cropped.) You will see a preview of the cropped image (Figure 2-11) with drag handles so that you can adjust the crop rectangle.



**Figure 2-11.** The 2.2 Crop tool in action

In GIMP 2.4, dragging any corner lets you resize the rectangle; dragging inside the rectangle lets you move it. In earlier versions of the GIMP, dragging the upper-left or lower-right drag handles lets you make the crop rectangle larger or smaller, while dragging the upper-right or lower-left drag handles moves the rectangle without changing its size.

Watch the mouse cursor carefully as you move it over the drag handles. The cursor will change as the mouse covers the handle, and will give you a clue as to what the GIMP will do if you begin dragging at that position. Figure 2-11 shows the GIMP 2.2 Crop tool: the cursor itself is a cross, plus a copy of the Crop tool’s “scalpel” symbol and a diagonal arrow indicating that you can change the rectangle’s size. The cursor for the other two corners is a pair of crossed arrows, which is used throughout the GIMP to indicate moving an object around within the window. GIMP 2.4’s Crop tool has simpler cursors.

At the same time, in GIMP 2.2 and earlier, the Crop dialog will appear. You’ll find precise information about the location of the upper-left corner, the width and height of the cropped image, and the aspect ratio. Or you can use the *From Selection* button, if you have already selected an area where you want to crop. In GIMP 2.4, there’s no more Crop dialog; the information shows up in tool options in the Toolbox. (If you find that the Crop dialog gets in your way in earlier versions, hold the Shift key down when you first drag in the image to start the crop. That will prevent the dialog from coming up.)



---

**Tip** If you already have an area selected, there's a faster way to crop: use the *Image* ► *Crop Image* item in the image window menus.

---

The dialog also offers a button for *Auto Shrink*, which attempts to shrink the crop rectangle to a hard boundary. On photographs, this usually won't do anything because there are no obvious "edges" to tell the GIMP which parts can be removed.

The Crop dialog's *X*, *Y*, *Width*, and *Height* fields can be useful for making small or very specific corrections. If you know the size you eventually want the image to be—for example, 640×480—you can simply type those numbers into *Width* and *Height*, and then use the Move handles to drag the rectangle to the right place. Or you can use the arrow buttons next to one of the fields to adjust that value by small increments.

You may find the arrow buttons confusing; in particular, the fact that clicking "up" on the *Y* position makes the crop rectangle go down, and vice versa. It happens because "up" on the *Y* position increases the *Y* coordinate, and pixel coordinates increase downward: the (0, 0) position is the upper-left corner of the image. If this confuses you, don't worry too much about it; just watch the crop rectangle when you make changes and notice in which direction it goes, and you'll be fine.

The arrow keys on your keyboard adjust the *X* and *Y* origin of the crop rectangle by single pixel values. The Control key combines with the arrow keys to change the width and height. Adding the Shift key to any of these combinations changes the values by a larger increment.

When you have the rectangle adjusted precisely to where you want it, hit Enter or click inside the rectangle to make it final. In 2.2 and earlier, you can also use the *Crop* button in the dialog.

The Crop tool's options are fairly simple, though the *Resize* option doesn't do what you probably expect it to. What it actually does is to reduce the size of the image canvas while leaving the layer size unchanged. These terms will become clearer when you read about layers in Chapter 3; for now, use *Crop* rather than *Resize*, and leave the other options at their defaults.

*Current Layer Only* restricts the crop to only the selected layer. You may find that useful once you begin to use multiple layers.

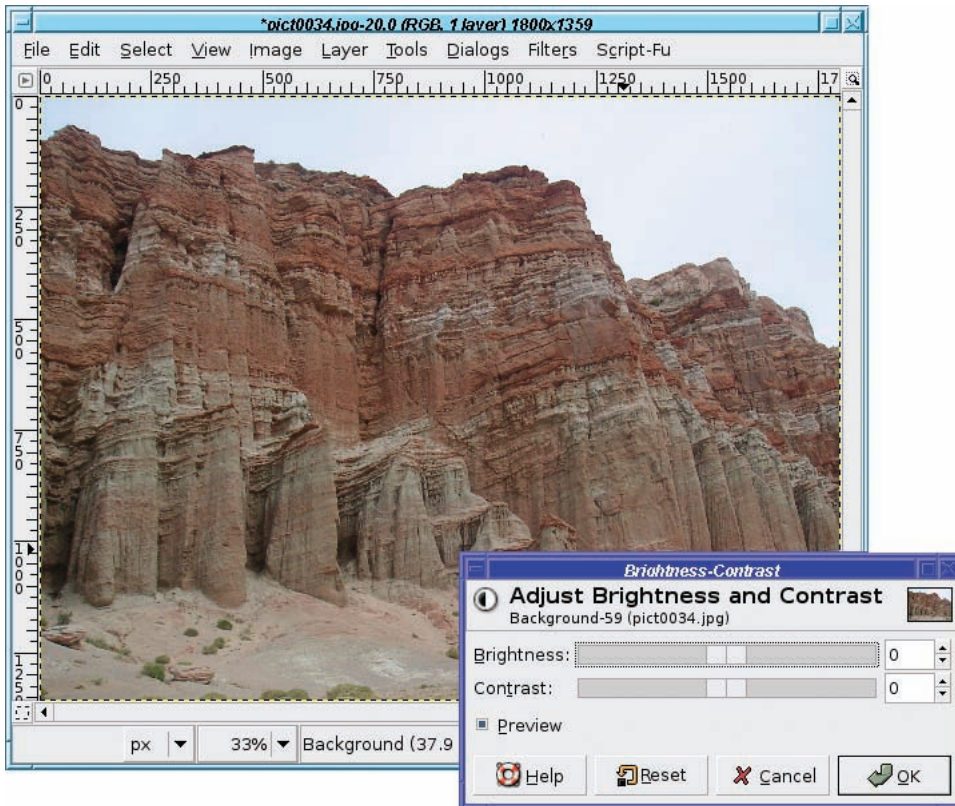
*Fixed Aspect Ratio* is the only tricky option. It only works once a crop rectangle is already active in the image. In order to crop an image and retain the image's existing aspect ratio, first click in the image, then when the Crop dialog appears, click on *From selection* to set the crop rectangle to the whole image. Then select *Fixed aspect ratio* in tool options. After all that, dragging on one of the resize handles will constrain the shape of the rectangle. With any luck this will become easier in future versions of the GIMP.

## Brightening and Darkening

A common problem with camera images is their being too light or too dark. The GIMP has three very powerful level adjustment tools to correct such problems: Brightness-Contrast, Levels, and Curves. You can find these tools in two different places in the menu system: the top-level *Colors* menu (*Layer* ► *Colors* in earlier versions of the GIMP), or *Tools* ► *Color Tools*.

### Brightness-Contrast

The simplest level adjustment tool is Brightness-Contrast (Figure 2-12).



**Figure 2-12.** *The Brightness-Contrast dialog. The original image is too dark.*

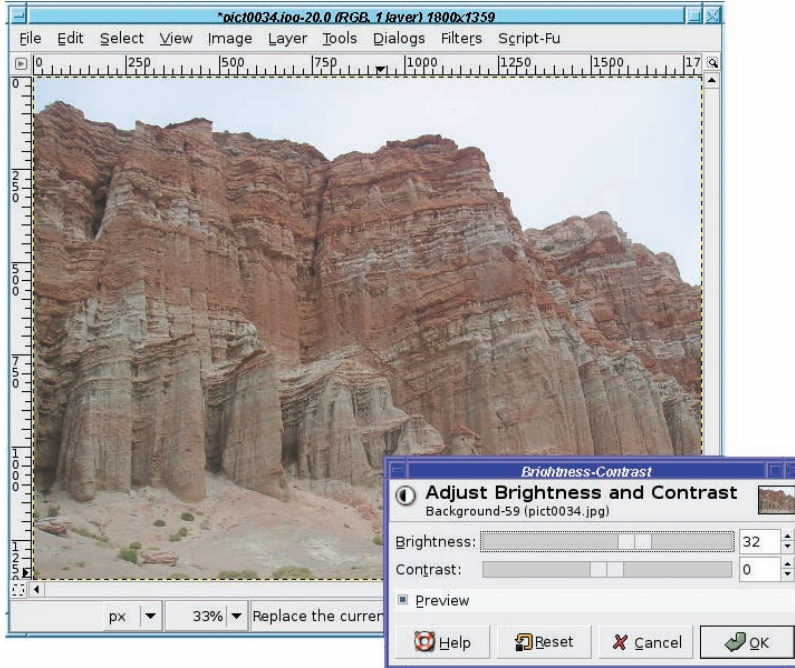
The basic controls in the dialog are simple: two sliders, one for brightness and one for contrast. You can also type in numbers (the values go from  $-127$  to  $127$ , with  $0$  as the current image), but most of the time you'll probably just drag the sliders until the image looks the way you want it.

In addition, the dialog has a checkbox marked *Preview*. With *Preview* on, any changes you make to the sliders will be shown in the image window, so you can see exactly how your changes will look. No actual changes will be made until you click *OK*. Many GIMP dialogs have a *Preview* option; it's a very useful feature that you'll want to enable unless you have a very slow machine.

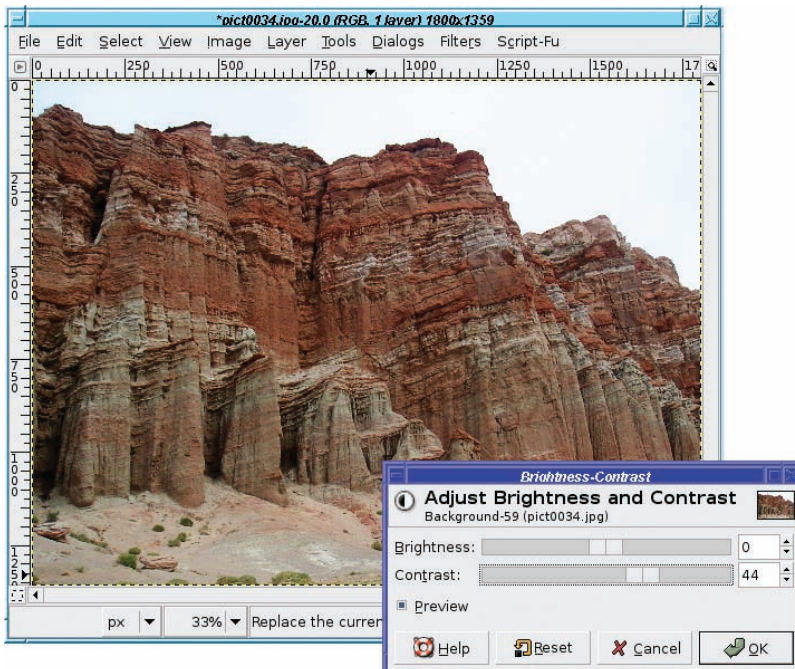
If you get the settings too far off, the *Reset* button at the bottom of the dialog will set them back to  $0$  so you can start over. *Cancel* will dismiss the dialog without making any changes to the image, while *OK* will take the current settings and change the image accordingly.

Figures 2-13, 2-14, and 2-15 show some examples of the changes brightness and contrast can make. Play around with the sliders to get a feel for what they do. You'll probably find that when you make an image brighter, you'll want to increase contrast a bit, or else it will look washed out. You may also be surprised at how well a slight contrast increase can bring out brighter colors in an image. Be careful of pushing contrast too far, though, or you can make the image look grainy.





**Figure 2-13.** Just brightness increased. The image looks washed out.



**Figure 2-14.** Just contrast increased. Notice how colors seem brighter, even though the image is still too dark.

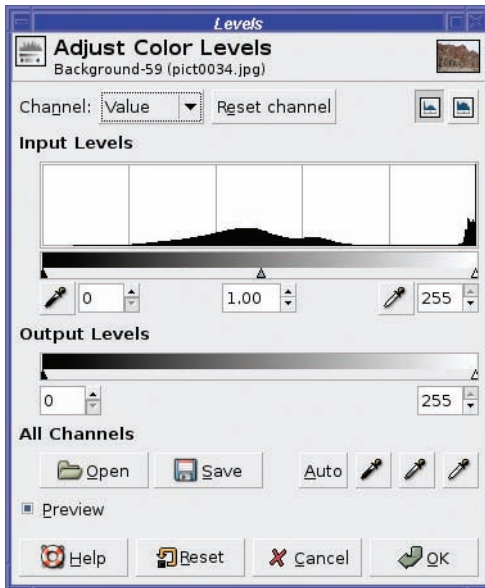


**Figure 2-15.** Both brightness and contrast have been increased.

## Levels

The Brightness-Contrast tool is straightforward and useful. But the GIMP has two other brightness tools that are much more powerful: Levels and Curves. As with Brightness-Contrast, make sure that *Preview* is checked so that you can see the effects as you play with settings.

Levels (Figure 2-16) has two sets of controls: *Input Levels* and *Output Levels*. *Input Levels* shows a histogram representing the current brightness, or *tonal range* of the image: this is a graph where the horizontal axis represents brightness, and the vertical axis represents the number of pixels that have that brightness. If the image is fairly well balanced, then the histogram should be spread fairly evenly across the image. Such an image is said to have *full tonal range*. If an image is way under- or over-exposed, the histogram is probably bunched up at one end or in the middle, or shows a spike over to the left or the right. Such an image has a *reduced tonal range*—it is not using all the possible pixel values that can be displayed.



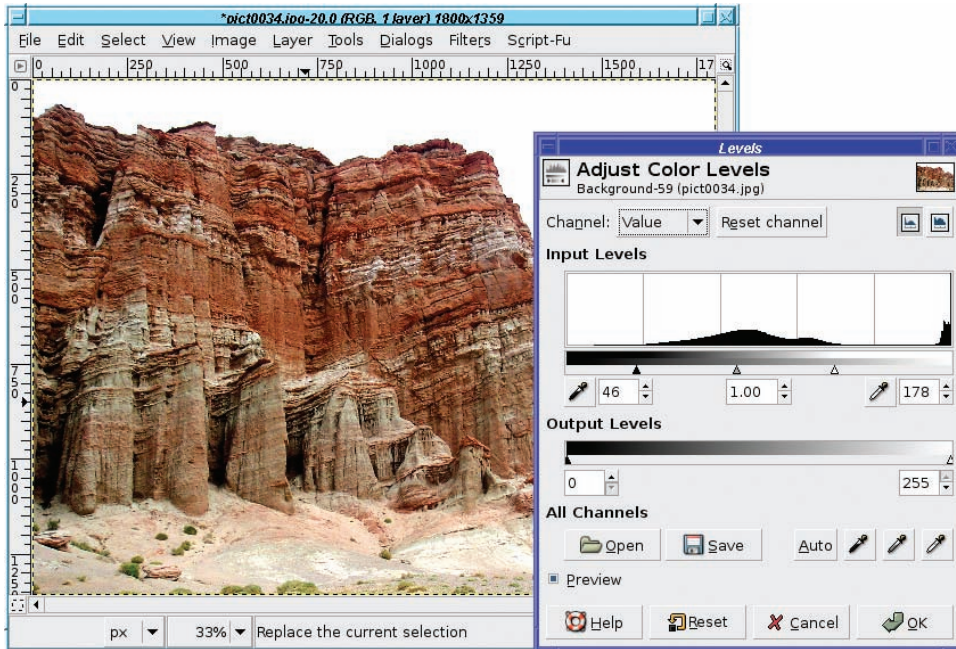
**Figure 2-16.** The Levels dialog. Notice the values bunched up at the right (light) end: that's the white sky in Figure 2-12.

Both input and output have sliders that you can play with to make the image brighter or darker, or give it more or less contrast. The three input sliders are referred to as *shadow*, *midtone*, and *highlight*, representing the darkest points in the image, a midpoint, and the brightest points.

Sometimes you can use the histogram for guidance: drag the left input slider to where it lines up with the left edge of the thick part of the histogram, and line up the right slider with the right side of the histogram (Figure 2-17). Don't neglect to experiment with the middle input level slider, which is also called *gamma value*—you can fine-tune the contrast of the image with it.

You can also use the *eyedropper* buttons under the input levels slider to pick points in your image that should be all black or all white: click on the left eyedropper button, and then click on the darkest point in the image. Then do the same to match the right eyedropper to the lightest point in the image.





**Figure 2-17.** Lining up the sliders with the histogram. The bits of the histogram bunched at the far right are because of the bright sky.

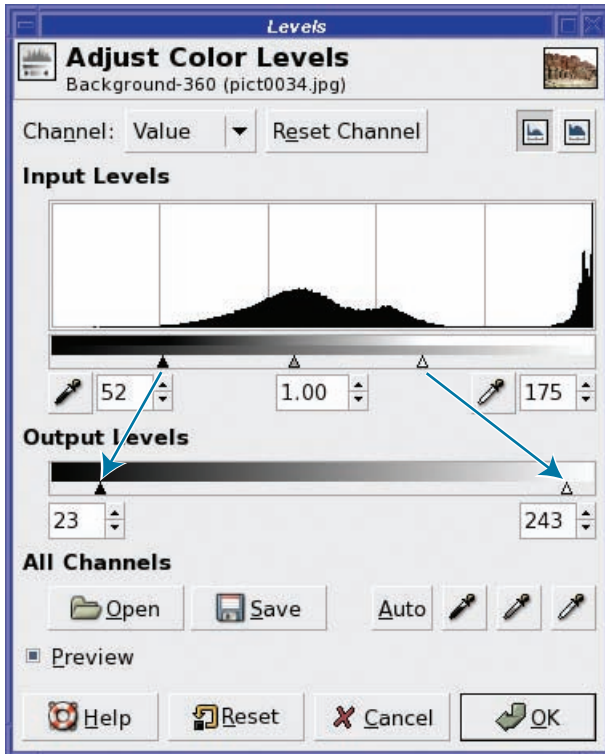
At the dialog's upper right are two buttons, *Linear* and *Logarithmic*. These control the way the histogram is drawn. They have no effect on the operation of Levels, but switching to *Logarithmic* can make the histogram easier to see by making it taller.

One common point of confusion is that the sliders on input and output seem to work in opposite directions. If you move the leftmost input slider (at the dark end) to the right, the image gets darker; if you move the leftmost output slider (also at the dark end) to the right, the image gets lighter. What's going on here, and how can you remember which slider does what?

Think about it this way: the input sliders let you expand the tonal range of an image that has a restricted range. The leftmost *shadow* slider controls the black level: anything in the original image that was darker than the position of this slider will be mapped to black. Similarly, anything to the right of (brighter than) the rightmost *highlight* slider will end up entirely white.

Once the image has been remapped according to the positions of the input sliders, the output sliders do the opposite: they let you restrict the tonal range of the final image. Moving the leftmost output slider right means that nothing in the final image will be fully black: anything that would have been black will now be a little bit less dark. Conversely, moving the rightmost output slider to the left means that anything that would have been white will now be a little less bright.

So input expands tonal range; then output restricts it. The leftmost input slider, which represents the blackest pixels in the image, gets mapped to wherever the black output slider is positioned; and likewise the right input slider gets mapped to the right output slider (Figure 2-18).




**Figure 2-18.** The input sliders, after the image's tonal range has been expanded, map to the output sliders to compress the final tonal range.

In most cases, full tonal range is a good thing; once you've adjusted the input sliders to expand your image's tonal range, you wouldn't want to restrict it again. Most of the time, you may not need the output sliders at all. They're there for the rare cases when an image really needs to have a restricted range, or for images that are legitimately lacking some part of their tonal range (no highlights, or no shadows).

## HISTOGRAM GAPS

If you open the Levels dialog on an image when you've previously made a brightness, levels, or curves adjustment, you may notice gaps in the new histogram. What do the gaps mean? The following image shows gaps in the histogram after correcting brightness. Logarithmic mode is used to make the histogram easier to see.



The gaps represent lost information. When the GIMP uses the input sliders to expand the tonal range of an image, some of the fine differences between pixels at the extremes of the image (in the very bright or very dark parts) may be lost. “Lost information” sounds bad, but most of the time, it wasn’t anything you cared about, or you would have set the sliders differently. But you can see the result in the histogram when you look at the image. This also means that you can sometimes use the GIMP as a forensic tool, to tell whether someone else’s image has been edited in this way.

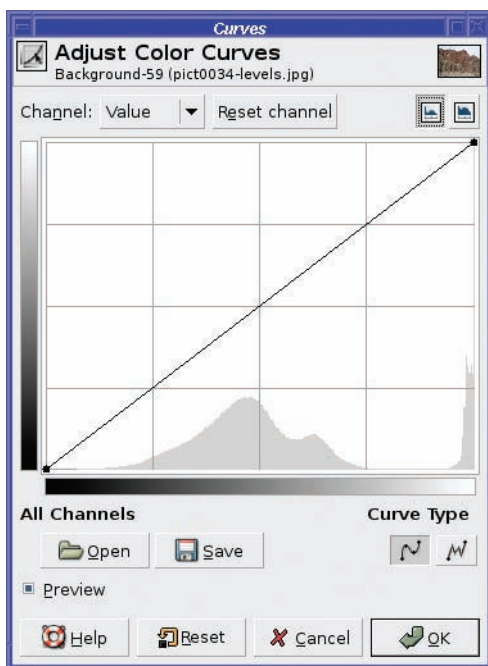
Another thing you may notice about the Levels dialog is the *Auto* button. This is theoretically comparable to the Auto Levels in some other photo-editing programs; but in practice, I find that it seldom helps photographs very much. (This may say more about which photos I try it on than it does about the GIMP’s Auto Levels.) It’s always worth a try; you can fine-tune from there if it doesn’t do quite the right thing, and you always have the *Reset* and *Cancel* buttons if you don’t like the effect.

## Color Correction

Finally, at the top of the Levels dialog is a menu allowing you to choose a *channel* (Curves has one, too). Initially this is set to *Value*, which you can think of as meaning “brightness”; but you can also choose *Red*, *Green*, *Blue*, or *Alpha*. The three color channels allow you to adjust color problems, such as a yellow cast from a photograph taken indoors without a flash. Color correction will be explored later, in the “Correcting Color Balance” section in Chapters 6 and throughout Chapter 8, but meanwhile, if you have an image that needs color correction, try using this menu in Layers or Curves. *Alpha* is a term meaning “transparency”; you’ll work more with transparency in Chapters 3 and 4.

## Curves

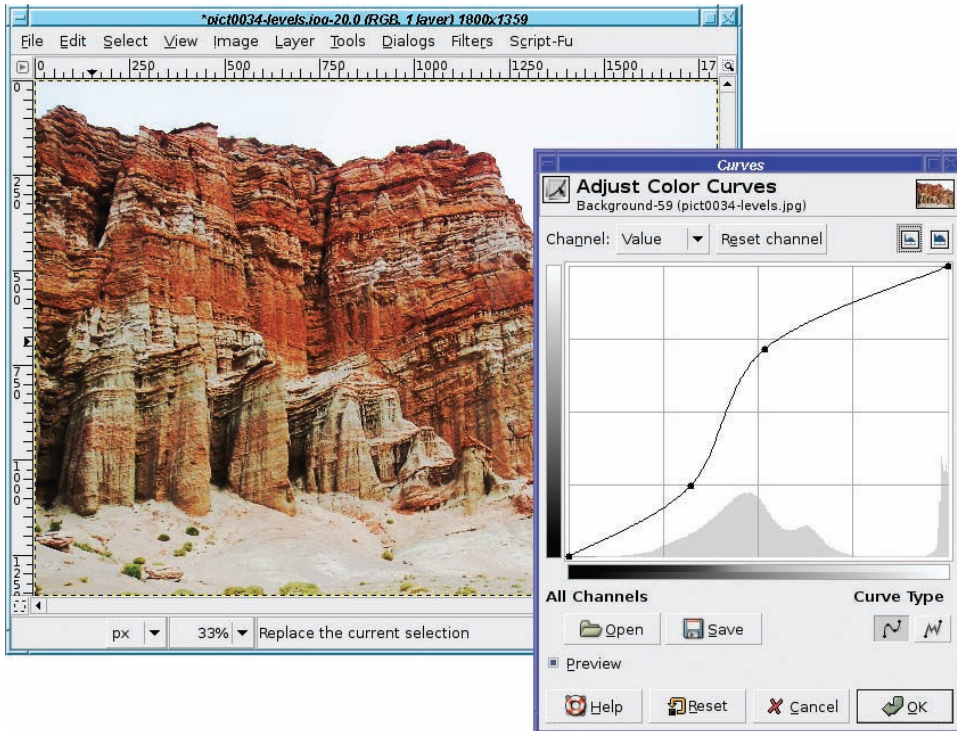
The Curves dialog (Figure 2-19) adjusts many of the same values as Levels, but in a different way. Some people find it more intuitive than Levels; others less so. You’ll probably find that you prefer one or the other; which one you use is purely a matter of preference.



**Figure 2-19.** *The Curves dialog*

In the Curves dialog, the horizontal bar at the bottom of the graph represents input values, just like the input graph in the Levels dialog. The vertical bar represents output levels. But you don’t need to know that to use Curves. To begin using it, just click on the line and drag it up a little (to make the image brighter) or down a little (to make it darker).

As soon as you click on the line, the GIMP creates a *control point*: a handle you can use to grab the curve and drag it higher (brighter) or lower (darker). You can go back and click again on an existing control point and slide it up or down along the curve, or you can leave existing control points where they are and click somewhere else on the line to create more control points (Figure 2-20). If you need to get rid of a control point you've made, just slide it all the way left or right to another control point or to one of the endpoints of the curve.



**Figure 2-20.** Control points in the Curves dialog

The advantage of the Curves tool (besides being fun to play around with) is that you can easily control which parts of the image are adjusted. The upper-right part of the curve corresponds to parts of the image that are already bright, while the lower-left part corresponds to darker areas. If most of the image is fine, but there are some bright areas that look blown out and you just want to tone those down, you can make a curve that's straight except at the upper-right, where it curves below the original line to make the bright areas darker. Then if the dark parts are too dark, you can grab the lower-left corner and drag it upward to brighten just those parts. You can make as many control points as you wish, and as complicated a curve as you desire, to fine-tune the correction you will make. (The catch is that such complex curves seldom accomplish much more than a single control point moved to the right place on the curve. But try it yourself. Experiment and see what you can do!)



## Options in the Curves Dialog

The *Linear* versus *Logarithmic* buttons at the upper-right of the dialog control the way the histogram is presented, just as in Levels; they have no effect on the operation of Curves.

At the lower-right are two buttons for *Curve Type*. By default, the Curves dialog will make a smooth curve connecting your control points; but if you want finer control than that, click on the *Freehand* button and draw your desired curve directly.

At the lower-left are buttons enabling you to *Save* a curve, or *Open* one you've previously saved. This may be helpful if you have a large number of images that all have basically the same exposure problems.

## Other Exposure Adjustments

In addition to the three general tools for adjusting exposure, the GIMP's *Colors* menu has several other useful tools.

Threshold (Figure 2-21) lets you map an image to black and white, adjusting the threshold point between the two. Like Levels and Curves, Threshold shows a histogram of the image's brightness. The area between the black-and-white sliders, colored blue, represents the range that will be white in the final image. You can drag the sliders directly, or mouse down in the histogram area and drag out the area that should be white.

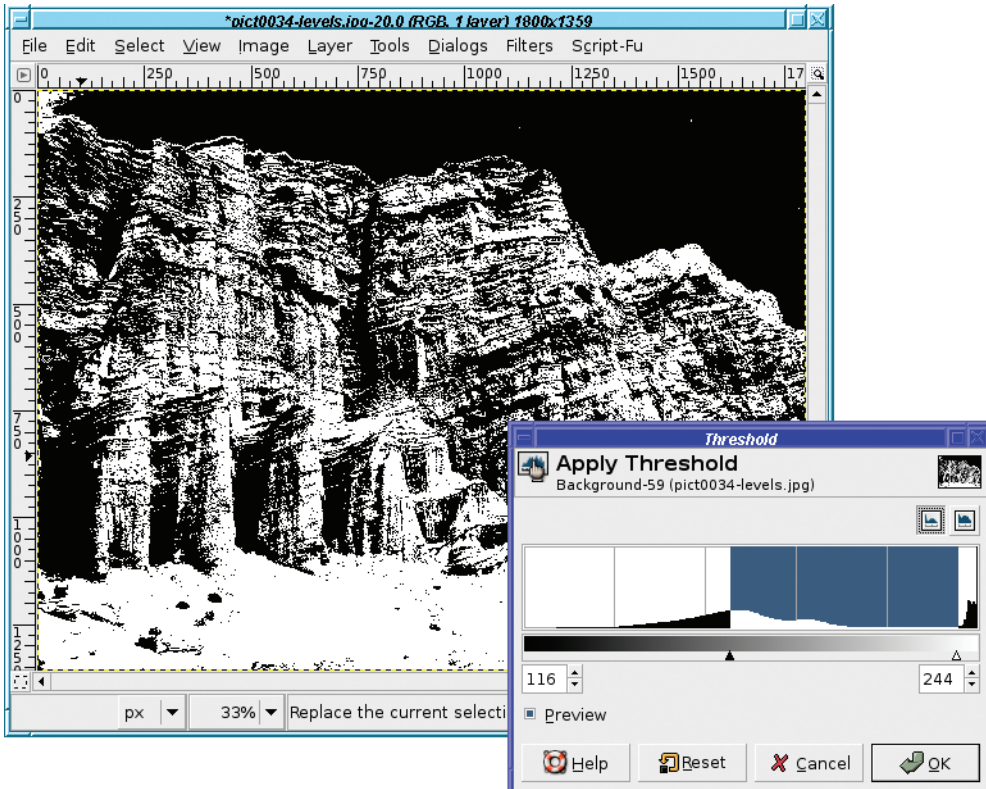


Figure 2-21. The Threshold dialog

**Tip** Threshold is particularly useful for scans of printed text documents and line art. If you end up with a few extra speckles of the wrong color, you can clean them up afterward using the GIMP's drawing tools (see Chapter 4).

Posterize (Figure 2-22) reduces the number of colors in the image, creating a gaudy display like you might see on an artistic poster. You can achieve a similar effect by converting an RGB image to indexed mode, but posterize is a quick way to achieve this effect, and it still leaves you with an RGB image in case you want to do further editing. The effect is most obvious when you use a very small number of colors.






Figure 2-22. A posterized image

Desaturate (Figure 2-23) removes the colors from an image, changing it to a grayscale image without requiring you to convert to grayscale mode. The three choices in the dialog let you choose the gray level three different ways, giving subtly different results. Desaturate has no preview, so you'll have to try it and use Undo if you don't like the results.



**Figure 2-23.** *A desaturated image*

Invert inverts every color in the image, turning it into its own photographic negative. This is most useful with black-and-white images, but can have an interesting effect on color images (Figure 2-24).



**Figure 2-24.** An image with the colors inverted

The *Auto* menu includes a set of automated functions useful for correcting photographs with poor exposure. They're worth trying if you have a difficult image you're trying to save; if you don't like the effect, there's always Undo.

- *Equalize* tries to spread the image's colors out among a wider range of intensities. Most of the time, you probably won't like the effect, but sometimes it brings out detail you didn't realize was there.
- *White Balance* tries to correct color casts, and often does a decent job. It's worth trying if you have an off-color image.
- *Color Enhance* makes the colors more intense—usually *too* intense, but try it and see.
- *Normalize* helps adjust the exposure on under-exposed images. On images that are well balanced except for being too dark, it may help immensely, but if the image's levels are uneven, it will probably just make things worse.
- *Stretch Contrast* and *Stretch HSV* are similar to *Normalize* except that they operate on the three color channels independently. Sometimes they can help remove color casts.



## Rotating

The GIMP offers two types of rotation. One is a quick fix for images shot in *vertical* format, while the other, free rotation, is a more subtle tool that can correct minor errors made while shooting.

### Rotating by Multiples of 90 Degrees

I shoot lots of *verticals*: photos where I rotate the camera by 90 degrees to the left or right in order to make an image that's taller than it is wide.

That's fine, but when I put it on a website, or load it into the GIMP, everything's sideways! The camera doesn't know that the image should be rotated, so it just saves it in the normal way, wider than it is tall. (Some cameras are smart enough to notice when they're rotated, and save the information to a file. The GIMP currently ignores this information, though some future version may read it.)

But fixing this in the GIMP is incredibly easy, using the *Image* ► *Transform* menu. You can *Flip Horizontally*, *Flip Vertically*, *Rotate 90 degrees CW* (that's clockwise), *Rotate 90 degrees CCW* (counter-clockwise), or *Rotate 180 degrees*. (The *Flips* are different from *Rotate 180 degrees* because *Flip* makes a mirror image, while *Rotate* preserves left and right if you look at the image upside-down.)


The final option in the menu, *Guillotine*, is a different operation (not a rotation) and will be discussed in Chapter 7.

### Free Rotation

Do you ever shoot a photo of a landscape, and notice later that you didn't hold the camera perfectly horizontal?

Okay, maybe *you* don't, but I have to confess I make that mistake from time to time. Fortunately, the GIMP comes to the rescue with the Rotate tool (Figure 2-25).

Select the Rotate tool in the Toolbox. There are several other tools that look similar to it, but remember, you can always hold the mouse over a tool (technically: *hover*) and read the tooltip to remind yourself which tool is which. Now look at the tool options—in particular, the *Transform Direction* option. For correcting a photo with a non-level horizon, set *Transform Direction* to *Backward*, and set *Preview* to *Grid*.




**Figure 2-25.** *The Rotate tool*

With the Rotate tool selected, click in the image to begin the rotation process. The Rotate dialog appears. You can use it to type in an explicit angle in degrees, or to change the center of rotation. Most of the time, though, you won't need to do that.


If you have *Preview* set to *Grid* or *Image+Grid*, you'll also see a grid of lines drawn on the image.

Now drag in the image and watch the grid change. The goal is to line the grid up to where the horizon *should* be in the image (Figure 2-26). When you're happy with how the grid lines up, clicking the *Rotate* button will rotate the image (Figure 2-27). Rotating a large image takes a while: GIMP's progress bar, at the lower-right of the image window, gives you an idea of how long the operation is taking.



**Figure 2-26.** *The Rotate tool in action: lining up the grid lines*

Notice that rotating makes the picture larger: the corners hang off the edges of the image. Behind the image, you see a gray checkerboard pattern. This pattern isn't really part of the image; it's how the GIMP represents transparency. Your rotated image is now on a transparent background. If you would prefer that the image be cropped to its original dimensions, with no transparency added, the checkbox for *Clip Result* in the Rotate tool options will do the crop for you. Alternately, you can increase the image size to show the corners using *Image* ► *Canvas Size*.



**Figure 2-27.** A rotated image (not clipped)

The Rotate tool has quite a few options.

- *Transform Direction* changes the direction of the rotation, and the *Preview* options work along with it. While *Backward* and *Grid* work perfectly for an image with a tilted horizon, there are other times when you want to rotate an object and see exactly how it will look. In this case, set *Transform Direction* to *Forward*, and *Preview* to *Image* or *Image+Grid*.
- *Affect* lets you apply the rotation to the whole image, or just one layer or path.
- *Interpolation* affects the quality of the rotated image. Unless your machine is very slow, you'll probably want to change this to *Cubic* or *Lanczos*.
- *Supersampling* improves the quality of the image still more, at the cost of being slower.
- *Clip result* will crop the rotated image to the same size as the original, as already noted.
- *Constraints* lets you restrict the rotation to be an even multiple of 15 degrees. If you don't check this option, you can still constrain the rotation by holding down the Control key while rotating.



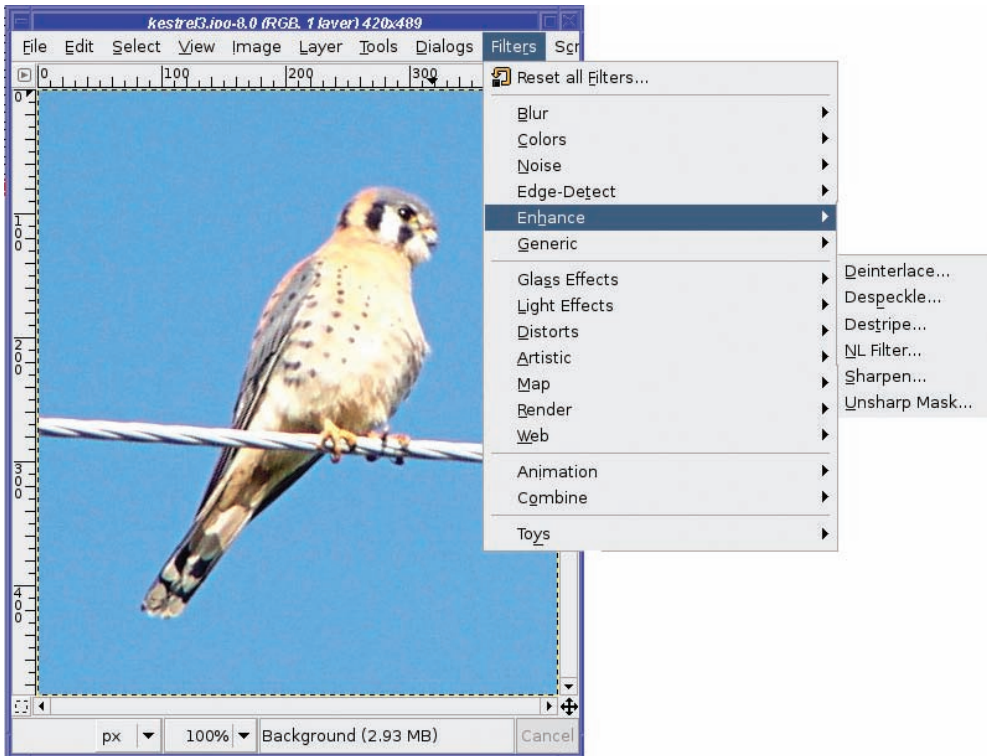
## Sharpening

Alas, our images aren't always as sharp as we'd like. (Blame it on the camera's autofocus! Yes, that's it!) For instance, the kestrel in Figure 2-28 came out a little out of focus.



**Figure 2-28.** A kestrel, out of focus. The GIMP can sharpen it!

The GIMP can't work miracles, but it can go a long way toward sharpening an image that's a little fuzzy. It offers several different sharpening methods. The two most important are Sharpen and Unsharp Mask, both located in the *Filters* ► *Enhance* submenu of the image window (Figure 2-29).



**Figure 2-29.** *The Enhance submenu*

Sharpen is very simple (Figure 2-30). It has a slider to change the amount of sharpening, and a preview window to show what the effect will look like. Scroll the preview window to the appropriate part of the image, or use Preferences to increase the preview size.



**Figure 2-30.** *The Sharpen dialog*

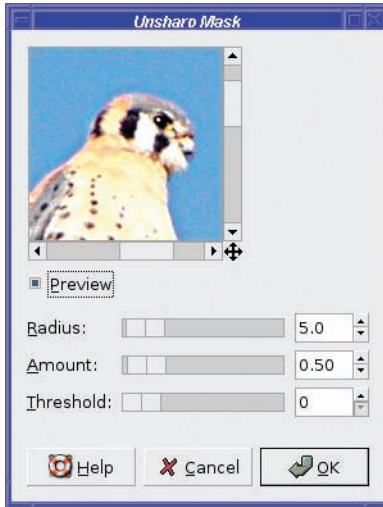
Here's what the kestrel looks like when sharpened (Figure 2-31).



**Figure 2-31.** *The kestrel after sharpening. The effects of Sharpen are usually fairly subtle.*

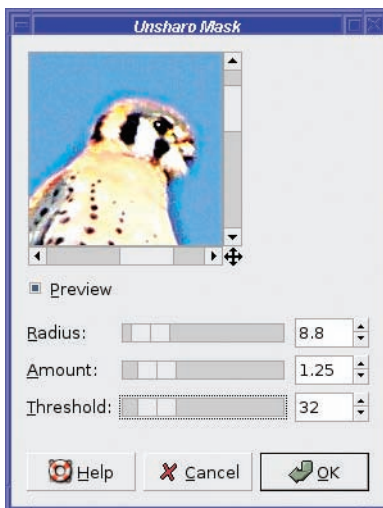
For more extreme sharpening, use Unsharp Mask (Figure 2-32). Don't be misled by the name, which comes from the details of how it works. Unsharp Mask makes an image much sharper—not less sharp—but it does that by making a blurry copy of the image, and then using that as a mask for the sharpening operation. Like Sharpen, it brings up a window with sliders and a preview area, but it has three parameters you can control, and it's capable of sharpening an image much more than the standard Sharpen.

Fiddle with the three parameters of Unsharp Mask and see what they do. Generally speaking, *Radius* affects the distance over which the sharpening works; make this bigger for large images, or very fuzzy ones. *Amount* controls how strong the effect will be; make this as small as you can so as not to over-sharpen and end up with a grainy image, or with halos around the object you're hoping to sharpen. *Threshold* controls the smoothness of the image; if you end up with grainy areas that should look smooth, such as the sky, an adjustment here may help.



**Figure 2-32.** *The Unsharp Mask dialog*

Unlike Sharpen, which can only create fairly subtle changes, Unsharp Mask can easily go too far in sharpening an image (for instance, see Figure 2-33). In fact, I find that the default values are usually already too extreme, and I decrease them slightly. Figure 2-33 shows the effect of the default values on the kestrel; Figure 2-34 shows what happens if I decrease *Amount* to 4, and increase *Threshold* to prevent the sky from becoming grainy.



**Figure 2-33.** *Too much sharpening makes an image look grainy or can give it a halo!*



**Figure 2-34.** Kestrel, sharpened with Unsharp Mask

## Fixing Red-eye

Flash pictures of people, especially children, often exhibit *red-eye*: pupils reflect too much light from the flash, turning the eye a diabolical crimson. This is particularly noticeable in photos of children (Figure 2-35) since their pupils are especially large.

There are lots of ways to fix red-eye, and the best ones require some advanced techniques. But here's a quick-and-dirty method that will work fairly well for most photos.

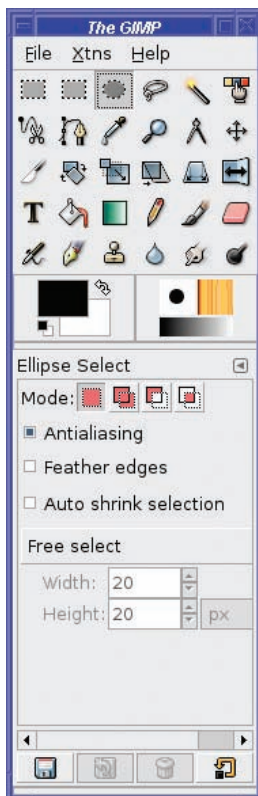
First, zoom in on the image, using the *View* ► *Zoom* menu, the *Zoom* button near the lower-left of the image window, or the + key. Zoom in a lot. Make the pupils nice and big. We'll change one pupil at a time.



**Figure 2-35.** *Red-eye can turn a sweet child into a demon!*

Now choose the Ellipse Select tool in the Toolbox (Figure 2-36).



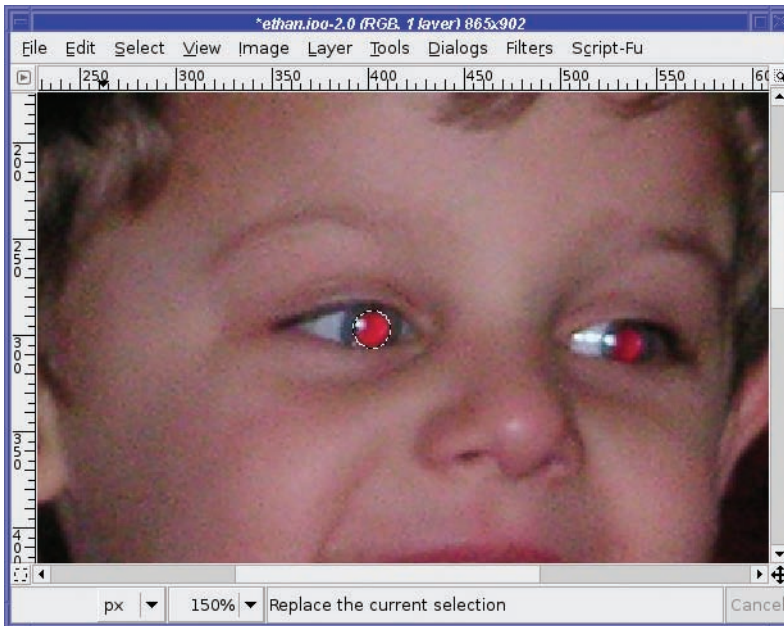


**Figure 2-36.** *The Ellipse Select tool*

Click slightly above and to the left of the pupil. Imagine a square box around the pupil—you want to hit the corner of where that square box would be. Then drag down and right (just like when you were making a crop rectangle with the Crop tool) until you have a circular selection around the pupil (Figure 2-37).

If it ends up in the wrong place, just undo and try again. It's not critical that it match the pupil exactly; just get it close. Unfortunately, the GIMP doesn't let you adjust the size of the ellipse interactively.

Now go to the *Colors* menu (in GIMP 2.2 and earlier, find it in *Layer* ► *Colors* or *Tools* ► *Color Tools*) and choose *Hue-Saturation*.



**Figure 2-37.** Make a circular selection around one pupil.

Without going too deeply into how colors are represented in the GIMP (see Chapter 8), the goal here will be to remove the red color from the pupil without losing the bright reflection.

First, click on the toggle button under *R*. This will restrict the operation to only the red color. (The eyes of animals sometimes show red-eye in colors other than red; if you're trying to fix a demon bunny or kitty, you may want to leave the setting on *Master*.) Then slide the *Saturation* control all the way to the left. *Desaturating* an image means removing all the color, or making it grayscale, though in this case, we're only removing all of the red.

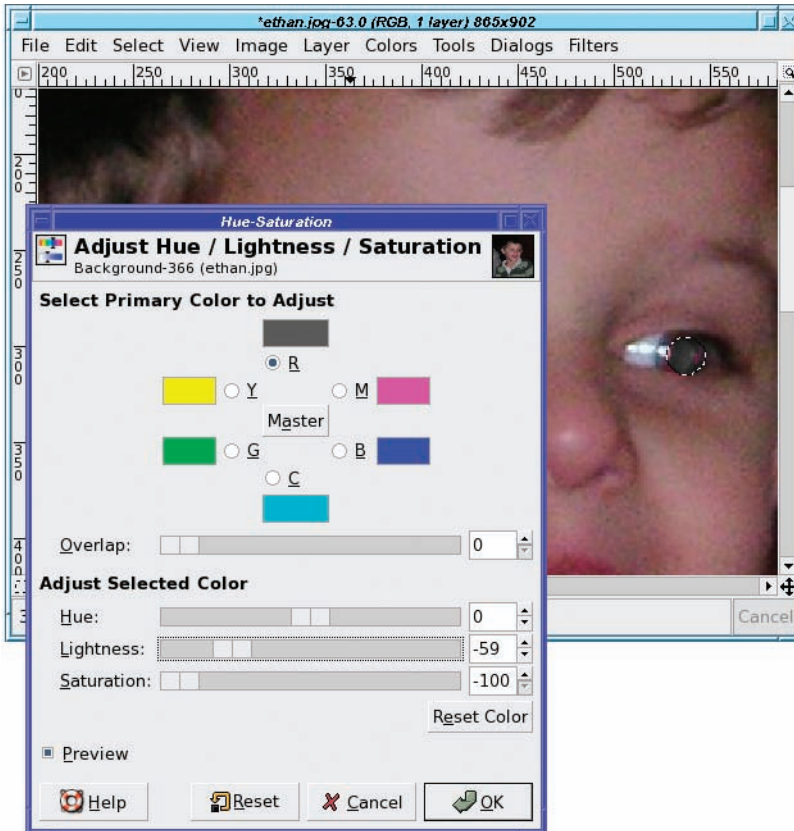
The hue-saturation operation only applies to what's selected, if there's a selection—otherwise it applies to the whole layer (you can think of it as applying to the whole image, until you begin working with layers in Chapter 3). That's true of most GIMP operations: if a selection is active, then only what's inside the selection will be changed.

With the pupil desaturated, the photo looks much better. The red is gone, and for some photos this may be all you need. But in this case, the pupil looks lighter than it should. I'd like to darken it a little.



If you find the selection boundary distracting, using Control+T toggles the selection outline off temporarily; it's a shortcut for *View* ► *Show Selection*. You can do this even when you're in the middle of an operation such as hue-saturation. Be sure to press Control+T again right afterward to turn it back on, or else you'll end up with a selection you can't see!

The *Lightness* slider makes the image brighter or darker, much as the *Brightness* slider in the Brightness-Contrast dialog did. With *Lightness* adjusted much darker (Figure 2-38), the pupil looks much more natural.



**Figure 2-38.** Making the pupil slightly darker

Now go back to the Ellipse Selection tool, select the other pupil, and apply the same operation. Voilà! No more demon! (See Figure 2-39.)



**Figure 2-39.** *No longer a demon!*

## Summary

Now you know enough to take your photos, fix any minor problems they may have, and share them with anyone in an appropriate size and format. You can crop your photos to preserve the most important parts and get rid of the rest. You can correct problems with brightness and contrast, minor rotation difficulties, and red-eye caused by too much flash. You know when to use GIF, when to use JPG, and when to use PNG or TIFF. You can even modify photographs that aren't completely in focus. That may be enough to keep you busy for quite a while.

But there's so much more that you can do with an image-editing program! The heart of image editing is learning to use layers. It's a different model than you may have used in simpler photo-editing programs, but once you begin to use layers, you'll wonder how you ever lived without them.

So take a deep breath, and prepare to explore layers, the real power of the GIMP.





# Introduction to Layers

**L**ayers sound intimidating. When you start to use an image-editing program, your first impulse is to avoid them and just make all your changes directly to the image.

But multiple layers make almost every aspect of image editing easier. If you learn about them now, you'll save an amazing amount of time, and operations that previously would have been impossible become easy. This chapter will cover the following topics:

- What is a layer?
- Using layers to add text
- Using the Move tool
- Changing colors: the GIMP's color chooser
- Using simple effects on layers
- Linking layers together
- Performing operations on layers
- Using layers for copy and paste
- A tour of the Layers dialog
- Bonus project: Making simple GIF animations

## What Is a Layer?


Every image the GIMP works with is made by combining one or more separate images, called *layers*, laid on top of each other.

As you'll recall from the discussion of file types in Chapter 2, images can include full or partial transparency. So can layers.

In the simplest way of combining layers, *normal mode*, anything completely opaque (not transparent at all) in the top layer is all you see in the final image. If the top layer is a photograph, like the ones you worked with in Chapter 2, then that's all you'll see in the final image, even if there are other layers underneath.

Layers get much more interesting when they include transparency. A transparent layer lets you put part of one image on top of another. Everywhere that the top layer is opaque is what you see in the final image; but where the top layer is transparent, you can see through to the next layer down. If that layer, in turn, has transparency, then you can see through to the next layer down, and so forth.

Figure 3-1 illustrates how a layer stack works in an image with several components. The top layer contains text; everything else is transparent. Next down is part of a flower photograph. The edges fade away just outside the flower so anything outside of that is transparent. Under the flower are layers for a picture frame, a shadow cast by the frame, and the background color. When these layers are combined, the image looks like Figure 3-2.



**Figure 3-1.** How multiple layers combine to form an image



**Figure 3-2.** *The image created by the layers in Figure 3-1*

Why bother with this? After all, you could just take a photo, put text on top of it, draw the frame directly on it, and get the same result. Right?

The difference is that with layers you have much more control. If you draw text directly onto an image, you'd better choose the right location for it, and the right font, size, and color. If you later decide the text would look better twice as big and over on the left instead of the right, there's no easy way to make the changes.

With layers, you can edit and revise to your heart's content. Set up a text layer. If you don't like it where it is, move it somewhere else. If you want to make it bigger, or move part of it behind another layer, or turn it off to see what the image looks like without it, you can do that too.

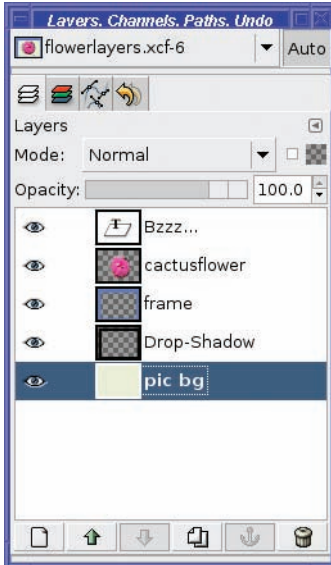
In addition, each layer contains only a specific part of the image. That can help when you apply other effects. A text layer contains only text, which means you can go back and change what it says, or try a different font or color. The frame in Figure 3-1 and Figure 3-2 was in its own layer, which made it easy to generate a drop shadow from it.

Since layers are so powerful, the GIMP is designed around them. Most of GIMP's effects (beyond the basic photo manipulation methods you learned in Chapter 2) either create new layers or assume that you're working from a single layer in an image that has other layers. Therefore, some familiarity with layers is crucial to getting the most out of the GIMP.

## Using Layers to Add Text

Layers are much easier to understand when you create them for yourself. So let's try a couple of basic projects.

The first step is making sure that the Layers dialog is visible. It's shown by default when you start the GIMP; but if you closed it at some point, bring it back now by using the Toolbox menu *File* ► *Dialogs* ► *Layers*. The Layers dialog looks like Figure 3-3 (the layers shown are the ones that correspond with the flower image above).




**Figure 3-3.** *The Layers dialog*

I'll talk about the various parts of the Layers dialog later in the chapter. For now, just make sure it's visible.


Now open an image—any image. I'll use this photo of a rock formation from Goblin Valley, Utah (Figure 3-4).

The Layers dialog now looks like Figure 3-5. The image only has one layer, named "Background".





**Figure 3-4.** A rock formation in Goblin Valley



**Figure 3-5.** The Layers dialog when the image is a single photo with no additional layers




---

**Note** Remember in the first paragraph when I said you could just make the changes to the image itself instead of using layers? Now you know it's really a layer anyway!

---

## The Text Tool

Now you'll add some text to the image. To do that, select the Text tool in the Toolbox (Figure 3-6).



**Figure 3-6.** *Selecting the Text tool*

I'll talk about the Text tool's options in a moment, but it's easier to see their effect when you have some text to work with. So click anywhere in the image. The Text Editor window pops up (Figure 3-7).




**Figure 3-7.** *The GIMP's Text Editor window, ready for you to type something*

The Text Editor is a very simple dialog:


- *Open* lets you load text from a file.
- *Clear* deletes whatever text you've typed in and lets you start over. Of course, normal editing operations also work, like selecting the text and typing something else, or using the Backspace key repeatedly.
- *LTR* and *RTL* stand for "Left To Right" and "Right To Left." Some languages require a right-to-left option, but for English text, you can safely ignore these buttons.

As soon as you type into the Text Editor, your text appears in the window. But something else happens: a new layer appears in the Layers dialog (Figure 3-8). You're now using layers!



**Figure 3-8.** *The new text layer in the Layers dialog*

I'll type some text in my sample (Figure 3-9). But it's tiny! My font size is much too small for this large image. Now it's time to look at the Text tool options.



**Figure 3-9.** *Font size is much too small*

Like all Toolbox tools, the Text tool has options shown below the Toolbox (or in the Tool Options dialog, if you've undocked it). You saw those options in Figure 3-6.

---

**Note** The title for this area is the name of the tool—in this case, *Text*—but this area, in the same place for all tools, is referred to as *Tool Options*.

---

You'll probably already be familiar with many of the Text tool options from using word processors. *Font*, *Size*, *Color*, and *Justify* all change the appearance of the text in the specified way.

*Hinting* and *Force auto-hinter* control how text is displayed at small sizes. In image editing, you'll be using fairly large text sizes most of the time, so you shouldn't need to change these settings. But if you're using small text, they're worth trying.

*Antialiasing*, however, is important to know. When it is on (the default), edges of letters will be smoothed by making jagged edges partially transparent. This is usually a good thing. However, in some cases you may not want this, such as in an indexed image with a very limited number of colors, or text in very small sizes where you want the letters to look sharp.

The three options *Indent*, *Line spacing*, and *Letter spacing* let you fine-tune the way characters are displayed. Most of the time you shouldn't need to change these, but some fonts may not correctly report their size (and therefore look awful) unless you adjust these numbers. Of course, sometimes you just want to space things out a little more than usual.

*Text along path* and *Create path from text* are two options that will become very handy once you begin using paths (Chapter 5).

If the text is too small, as in this case, the solution is to adjust *Size*. You can click on the up and down arrow buttons, or type a different number into the text field.

You may also notice that *Size* has a units menu next to it. By default, the GIMP uses pixel sizes for fonts, but you can specify a different unit, such as points, inches, or millimeters.

The text will change in the image window as soon as you change any of the Text tool options.

Don't count on the Text tool creating text in exactly the size the units menu leads you to expect. The size at which text is drawn on an image depends on many factors, such as the size of the image, the current resolution setting, and options such as hinting. Choose a text size that looks right in the image, rather than trying to force it to a specific number of points or inches.

With the text size set to 150 pixels, the text looks much better (Figure 3-10).

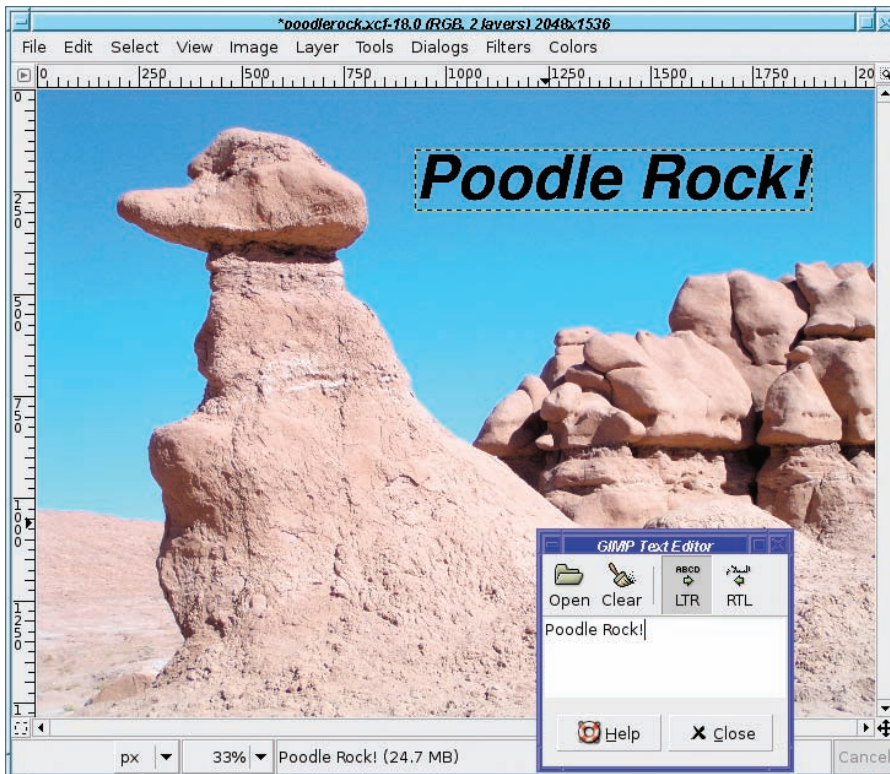



Figure 3-10. The font size is all better now.

## Using the Move Tool

Since the text is in a different layer, you can move it to different places in the image, and see where you like it best.


For this, you need the Move tool (Figure 3-11). The Move tool's icon is a set of four arrows, pointing up, down, left, and right. Become familiar with this glyph: you'll see it used to mean “move” in many places throughout the GIMP.



**Figure 3-11.** *The Move tool*

For instance, when the Move tool is active, moving your mouse over the text in the image window changes the cursor from a pointing hand to an arrow with the Move tool's crossed arrows next to it: the Move cursor (Figure 3-12).

There's another, less obvious way of getting to the Move tool: press and hold down the spacebar on your keyboard. As long as the spacebar is pressed, the Move tool is active, no matter which tool was previously selected in the Toolbox. When you release the spacebar, the GIMP will go back to the previous tool. This is an especially useful shortcut for moving newly created text layers or newly pasted layers.



**Figure 3-12.** At left is the cursor when not pointing at the active layer. At right is the Move cursor when pointing at the active layer. The red arrow points to each cursor’s “hot spot.”

---

**Tip** Every cursor shown in a computer program has a *hot spot*, the location where the cursor points. In the case of the Move cursor, the tip of the arrow is the important part and the rest is just decoration to let you know which layer will be moved. For the Hand cursor, the tip of the pointing finger is the hot spot.

---

Why does the cursor change? The answer lies in the Move tool’s single option: a choice between *Pick a layer or guide* and *Move the current layer*.

Every image in the GIMP includes the notion of a currently selected layer, or *active layer*. This is the layer that is highlighted in the Layers dialog (and that’s why it’s helpful to keep the Layers dialog visible). Back in Figure 3-8, the active layer was the text layer that had just been created.

You can change the active layer at any time by clicking a different line in the Layers dialog (click either on the layer name or the preview).

By default, the GIMP’s Move tool uses *Pick a layer or guide*. In this mode, if you click somewhere in the image, whichever layer is visible at that location will become the active layer and will move if you drag it.

The Move cursor, with the crossed arrows, shows that you’re pointing at the active layer. If you click and drag, that’s the layer that will move.

If your mouse is outside the current layer (in this case, that means anywhere that it’s not directly over text), the Move tool will switch to a different cursor to warn you that you might not be moving the layer you expect.

---


**Caution** In GIMP 2.2, *Pick a layer or guide* only makes the chosen layer active while it is being moved. After you release the mouse button, whichever layer was previously active becomes active again. In GIMP 2.4, the layer stays active even after you release the mouse button.

---

Go ahead and experiment (Figure 3-13). Try moving your mouse over the text while watching the cursor. Make sure it shows the crossed arrows indicating “move.” Then drag the text layer to a new location. Also, try dragging when you’re not over the text layer (no Move cursor showing). Notice that in this case, the background layer—not the text layer—is the one that moves.

Remember, Control+Z or *Edit* ► *Undo* will undo any layer movement you don’t want to keep.





**Figure 3-13.** Find the right spot for your text.

When moving text (especially tiny stuff), it can sometimes be tricky to get your mouse over the text and keep it there. The cursor has to be exactly over part of a letter, not just near it.

In the Move tool's alternate mode, *Move the current layer*, the active layer will always be moved when you drag *anywhere* in the image. This is safer—you don't have to be as careful where you put your mouse, or watch the cursor carefully to see where it changes. However, it can be inconvenient if you're moving several layers, since you have to go back to the Layers dialog and switch to a new active layer for each move.

Regardless of which Move tool mode you decide to use, you can switch to the other mode temporarily by pressing the Shift key before you click. Try both modes, and see which one you prefer.

When the Move tool is selected in the Toolbox, pressing the arrow keys (up, down, left, and right) will move the active layer one pixel in the indicated direction. Combining the Shift key with one of the arrow keys will move the active layer by 25 pixels instead of just one.

## SAVE YOUR WORK

To save the results of your hard work when you're gimping an image with multiple layers, use the GIMP's native format: XCF.

The XCF format saves everything about the image: all the layers, the current selection, any paths or guides (those will be discussed in later chapters)—in short, all the details you need to continue working on the image later.

In the GIMP's Save As dialog, if you type a file name ending in ".xcf", the GIMP will automatically save to the right type. Of course, once you save as XCF, it will continue doing so until you tell it otherwise.

Since XCF files can be quite large, you can use .xcf.gz or .xcf.bz2 to tell the GIMP to use the gzip or bunzip compression formats.

The disadvantage of XCF format? No other program can read it. You should definitely use XCF to save a copy of the image with all the parts preserved. But once you've finished editing, you'll want to use *File ► Save a Copy...* to save another copy of the file in a more standard format, such as JPEG, PNG, or GIF (see Chapter 2, "Image File Types").

Save frequently. Get into the habit of pressing Control+S whenever you get to a place where you're satisfied with a set of changes. Then you're protected against unforeseen events such as computer crashes, power losses, or alien invasions.

## Changing Colors

Now that the text is the right size, it becomes more obvious that it's boring in plain old black. It's time for another color.

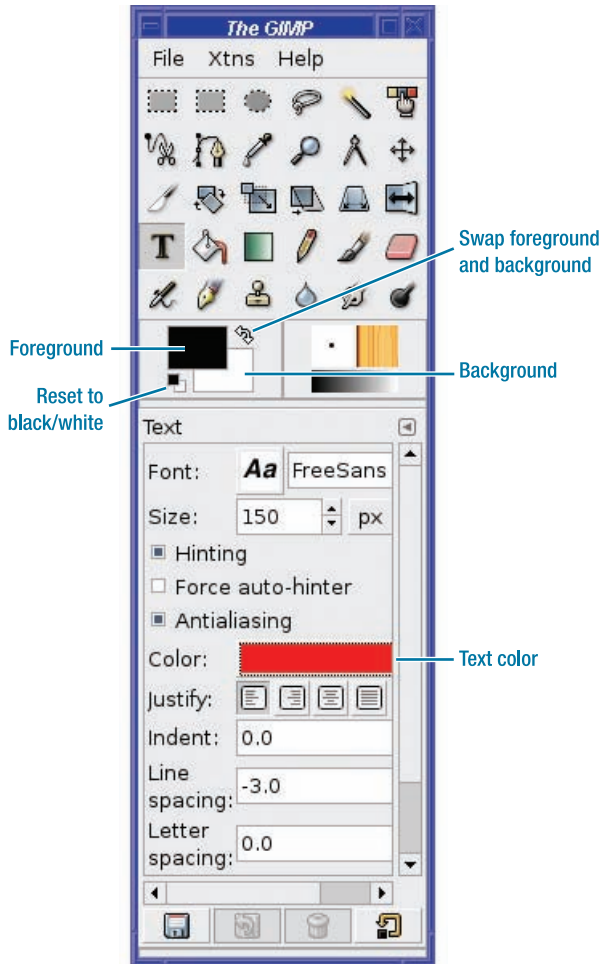
When you created text, it showed up as black because it used the GIMP's current foreground color, the foremost of the two color swatches in the Toolbox window (Figure 3-14). (It's black in Figure 3-14, meaning black is the current foreground color.) Any text you add to the image, or any lines you draw, will be in this color.

Behind the foreground color is another swatch showing the background color. The most common use of the background color swatch is to save a second foreground color. You can swap foreground and background by clicking on the small arrow to the upper-right of the two swatches.

The background color does have other uses more consistent with its name. Erasing or clearing uses the background color in images that lack transparency. Some drawing tools have options to use the background color instead of foreground in certain circumstances (you will see an example at the end of this chapter and more details in Chapter 4).

To change either the foreground or background color, click on the appropriate color swatch in the Toolbox.





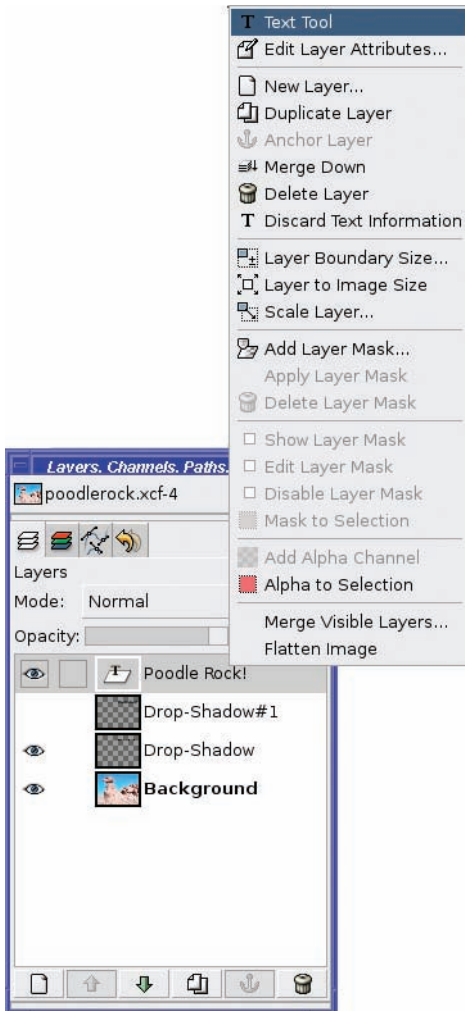
**Figure 3-14.** Color controls in the Toolbox window

Sometimes you may have to click twice. If the first click doesn't do anything, try again.

The tiny black-and-white boxes to the lower-left of the swatches are a quick way of resetting the colors to black and white.

Figure 3-14 also shows the Text tool's color as different from the foreground color in the Toolbox. To change the color of an existing text layer, with the Text Editor window showing, click on the bar next to *Color* in the Text tool options to bring up the color chooser. Like the color swatches, this changes to mirror the color of the text layer.

If you need to change text, color, or other attributes of an existing text layer that is no longer being edited, you can do that from the Layers dialog. Right-click on the text layer, and choose the first menu item: *Text Tool* (Figure 3-15). This will bring back the Text Editor window. If you don't need the Text Editor and merely need to change color, font, size, or other tool options, try selecting the layer in the Layers dialog, then clicking on the text in the image window. (A second click on the text will bring back the Text Editor window.)

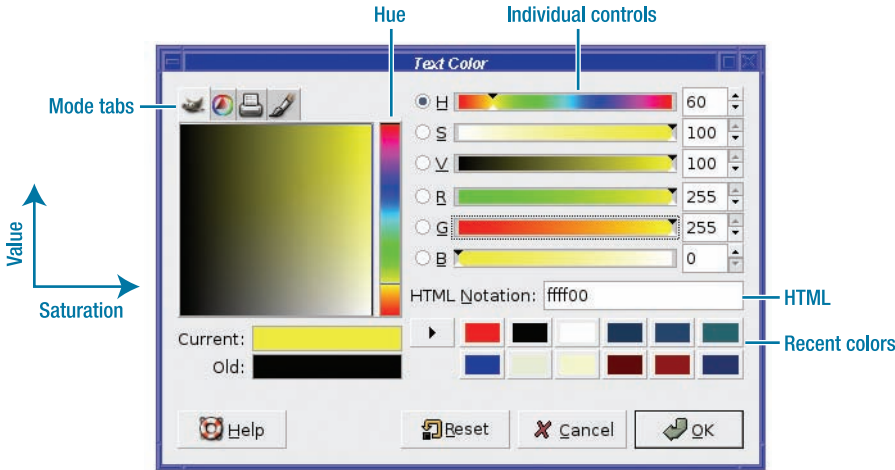


**Figure 3-15.** *The Text Tool item in a layer's context menu*

How do you know when the tool options are telling you about an existing text layer, instead of just the default settings for a new layer you might create? Watch the tool options to see if they change to the values you expect for that text layer. If it's still hard to tell, check the *Text along path* and *Create path from text* buttons at the bottom of the Text tool options: they're grayed out for new text, but active for existing text. I'll talk about paths later, in Chapter 5.

## The GIMP's Color Chooser

The color swatches in the Toolbox and the Color button in the Text tool options both bring up the GIMP's color chooser (Figure 3-16).



**Figure 3-16.** *The Color Chooser dialog*

Don't panic! Choosing colors isn't hard, even if the dialog looks complicated.

Let's start with the easiest way to choose a color. First pick a hue, using the rainbow-colored vertical slider. Click near the shade you want. If you "feel lucky" you can just stop there, or you can modify the hue by dragging the slider up or down. (Unfortunately, the arrow keys don't help here.) Then, use the square area to the left to adjust the brightness (value) and intensity (saturation) of the color.

Try to click in a place that shows a color close to your goal. Then slide around, watching the *Current* sample area underneath.

To get a bright color like the ones shown in the hue slider, click in the value-saturation square and drag all the way to the upper right. Dragging all the way to the left will always give black, regardless of the hue; dragging to the lower-right will always give white. (For a discussion of HSV versus RGB color, see the first two sections of Chapter 8.)

When you see the color you want, stop!

### OTHER COLOR ADJUSTMENT MODES

Other ways of adjusting the color in this dialog include dragging the individual controls for Hue, Saturation, and Value (HSV) or for Red, Green, and Blue (RGB), or clicking on one of the recent colors buttons, if one happens to show a color you want. If you have an HTML page that specifies a color by number, like #0276f4, you can match it by typing the number (without the #) into the *HTML Notation* field. If you then hit a tab, or click in some other field to indicate that you're finished typing, you'll see the new color.

There are lots of other ways to select color in the GIMP. Try clicking on the mode tabs to see some of the other types of color choosers the GIMP offers. Chapter 8 will discuss some other methods.

The *OK* button finalizes the color and applies it to the appropriate color swatch. Or, if you got to the color chooser by clicking on the Text tool's color button, the new color will be applied to the current text layer.

Figure 3-17 shows the text, now yellow.



**Figure 3-17.** *The text has been changed to yellow.*

## Simple Effects Using Layers

I mentioned earlier that a lot of the GIMP's special effects either assume that you're using layers, or they create new layers of their own.

A good example is the *drop shadow*.

### Drop Shadows

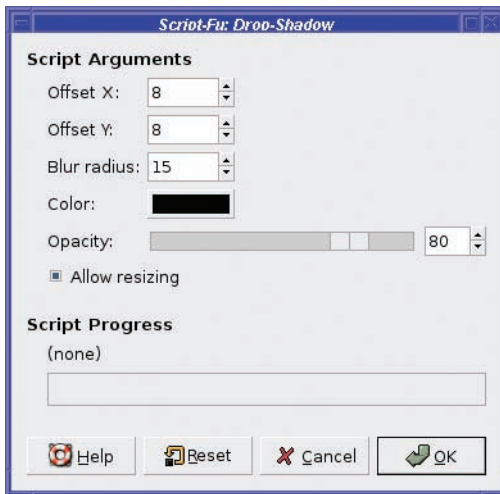
Text on its own, in a single color, can look flat and boring. Adding a shadow makes it look more three-dimensional, like it's floating above the page. *You* know your text is on a layer by itself, but a drop shadow makes it look like a separate layer to everyone else!

A drop shadow has another useful property: it adds a dark edge to light-colored text (or vice versa). This helps if you need to show text against a complex background that includes both light and dark colors.

A drop shadow can only be generated under any layer with transparent parts, such as a text layer. Transparency is how the GIMP decides where the shadow will be.

To make a drop shadow for a layer, make sure the right layer is highlighted in the Layers dialog, and then go to the menu system:

*Filters* ► *Light and Shadow* ► *Drop-Shadow...* (previous versions of the GIMP will show it in *Script-Fu* ► *Shadow* ► *Drop-Shadow...*). This brings up the Drop-Shadow dialog (Figure 3-18).



**Figure 3-18.** *The Drop-Shadow dialog*

*Offset X* and *Y* represent how far (in pixels) the shadow will “drop” from the layer that casts it. Larger offsets make the object appear to be floating higher off the page; but increase the offset too far and the shadow won’t look realistic.

If you decide later that you want a larger or smaller offset, you can use the Move tool to drag the shadow layer toward or away from the text layer.

*Blur radius* controls how blurry the shadow’s edges will be. If you increase the offset, the shadow may not look right unless you increase the blur as well.

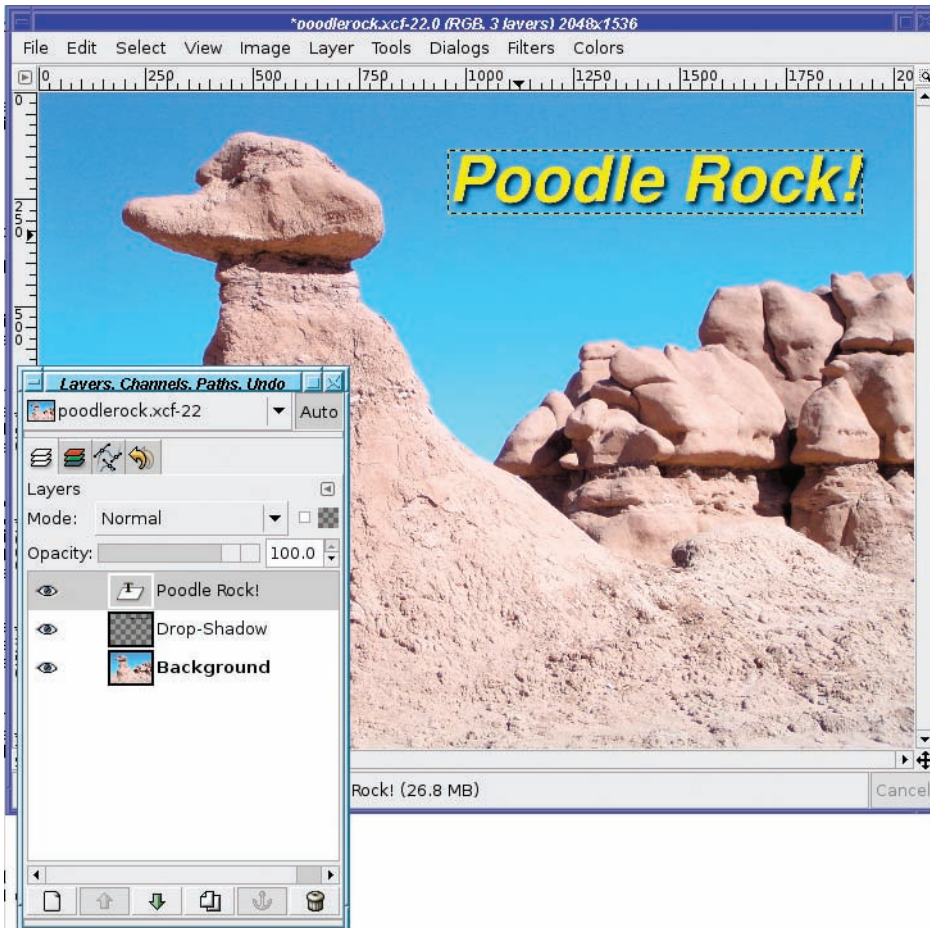
Hold your hand above a table and observe its shadow. Notice how it gets sharper as you move your hand closer to the table, and blurrier as you move your hand higher. Our brains expect shadows to look this way. If you want a realistic shadow, it’s best to try to match this.

*Color* controls (surprise!) the color of the shadow. Usually you want a black or dark gray shadow, but tinting it slightly (using a very dark blue, for instance) can sometimes add an impression of richness. You can also make drop shadows with bright colors, or with white: they won’t look like shadows, but they can help make dark text stand out against a dark background.

*Opacity* controls whether the shadow will be opaque or partially transparent. With a real shadow you can often see the pattern underneath. Less opacity creates a similar effect by letting you see a bit of the layer below. The opacity number is a percentage, from 0 (fully transparent) to 100 (fully opaque).

The *Script Progress* area is where GIMP will show a progress bar while the script runs.

Figure 3-19 shows the effect of drop shadow with the default settings. Notice that in the Layers dialog, a new layer named “Drop-Shadow” has been added for the shadow, but the text layer is still the active layer.



**Figure 3-19.** Text with drop shadow added

The first time you make a drop shadow, I recommend using the default settings. If you decide you need an adjustment, you can always undo and try again with different settings.

---

**Tip** This is one of those times when using the *tear-off* feature of the context menus can really pay. You can keep the *Light and Shadow* menu up and click *Drop-Shadow...* over and over as you adjust parameters.

---



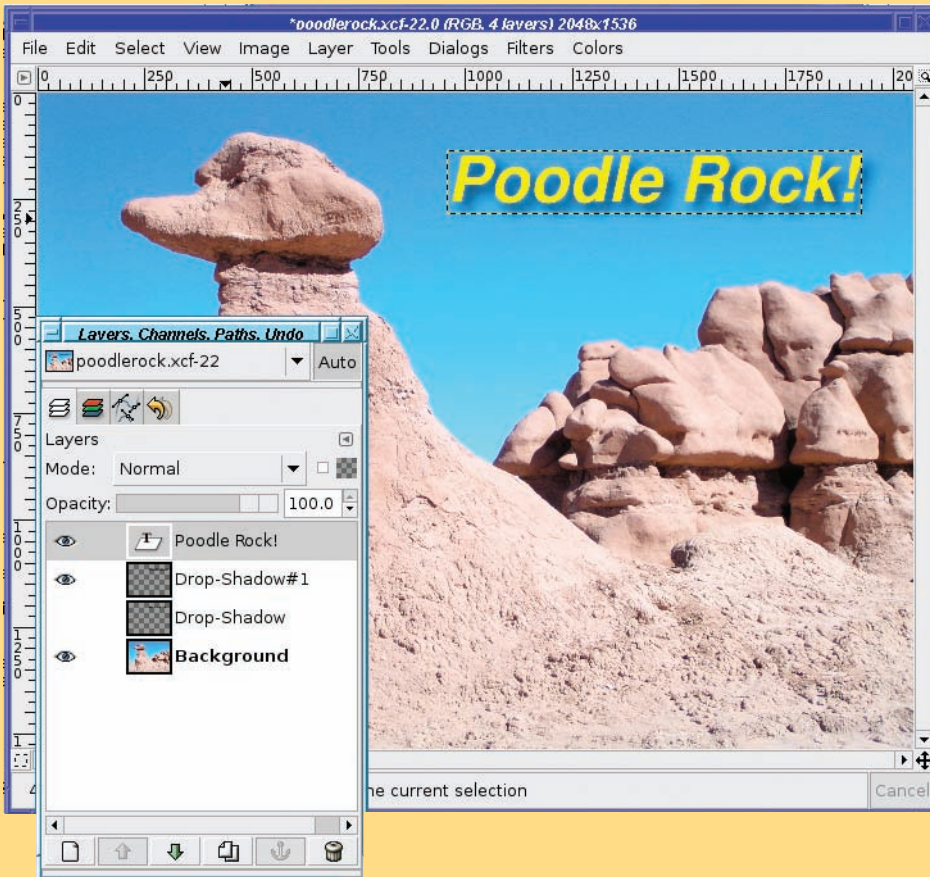
## LAYER VISIBILITY

Wait! Don't undo your first drop shadow just yet. The Layers dialog has another nifty feature: you can turn layers on and off to try alternate versions of the same effect.

Notice the “eyeball” icon next to the left of each layer in the Layers dialog? This is a button controlling visibility of the layer. Click on the eyeball next to the “Drop-Shadow” layer, and the shadow vanishes. You haven't deleted it from the image, but it has temporarily been made invisible.

When you save as XCF, the invisible layers will also be saved, just like other layers.

If you make another drop shadow with different settings, the GIMP will make a new layer, this time named “Drop-Shadow#1”, as shown in the following image:



With both drop shadows stored as layers in the image, it's easy to compare them. Make first one, then the other visible. You can even try making them both visible at once, though with drop shadows that usually won't work very well.



Beyond drop shadows, you'll find some other interesting effects you can try on text layers in the *Filters* ► *Alpha to Logo* submenu (in earlier versions, *Script-Fu* ► *Alpha to Logo*). Each one starts with an existing layer (most often they're used for text, but they don't have to be) and transforms it in some whizzy way. Try them all!

---

**Tip** This is another place where using tear-off menus can really help.

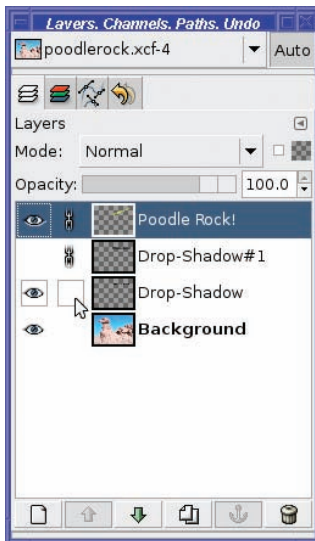
---

## Linking Layers Together

Once you've run a text effect or generated a drop shadow, you're left with several layers that have very specific relationships. If you drag your text layer with the Move tool and it leaves the shadow behind, you've lost your carefully calculated shadow offset.

Fortunately, the GIMP offers a solution: you can link several layers together.

In the Layers dialog, the space between the “eyeball” visibility icon and the layer preview is a *chain link* (Figure 3-20). Initially it's a blank space, but clicking in this space makes the chain link icon appear.



**Figure 3-20.** *Chaining layers together*

When you turn on chain link icons for several layers, they will be linked together. (It doesn't make sense to link just one layer. What's the sound of one hand clapping?) Select any one of the layers and drag with the Move tool, and they will all move together.

Linking works on invisible layers, too. If you have several alternate drop shadows, as in Figure 3-20, go ahead and link them all. They'll move together, and if you later decide to use one of the alternatives, it'll be in the right place.

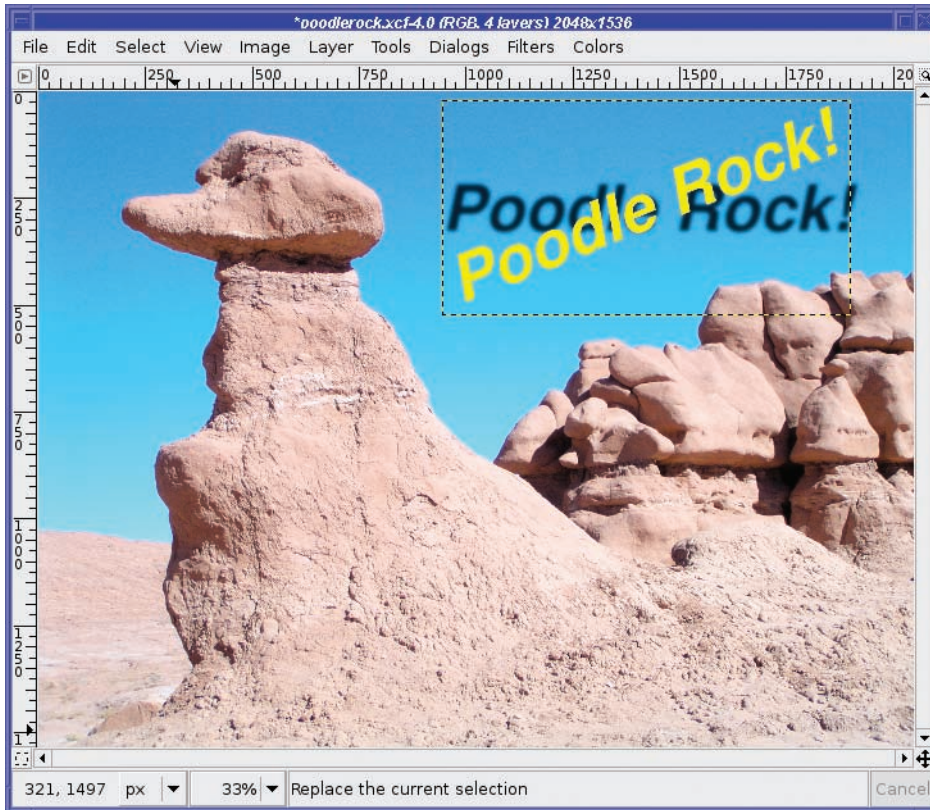
Other operations are affected by layer chaining as well, as you'll see in a moment.

## Performing Operations on Layers

Most GIMP operations work only on the selected layer, not the whole image. The drop shadow is a fine example.

The Rotate tool you used in Chapter 2 actually works on layers, not the whole image. You didn't notice because that image *had* only one layer. But when you have multiple layers, you can use that to your advantage.

For instance, you can rotate just a text layer (Figure 3-21).



**Figure 3-21.** *Just the text layer rotated*

Whoops! What happened here?

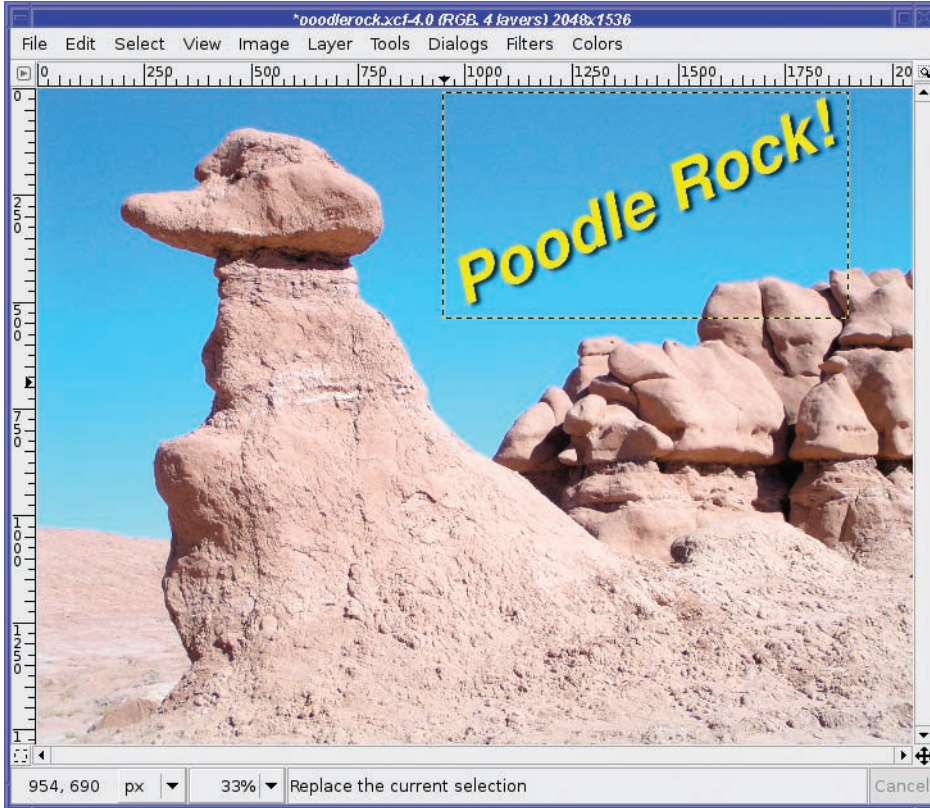
In this case, the GIMP rotated the text layer—but the drop shadow layer is still in the same place. That doesn't look good at all.

There are a couple of ways you could solve this.

You could rotate the text layer first, and then create the drop shadow later. That works, but it's not helpful when you've already spent time making a drop shadow that looks just right.

Or you could make a note of the number of degrees you rotated the text, and then apply the same rotation to the drop shadow layer. Ick! Computers are supposed to handle repetitious tasks for us.

Fortunately, if you chain several layers together, the Rotate tool will treat them as one. The preview will only show the selected layer, but the final rotation will include the whole chain (Figure 3-22).



**Figure 3-22.** Text and drop shadow rotated together

Unfortunately, chaining layers together doesn't work for all GIMP operations. It works for most of the "transformation tools," like Rotate, which distort the image in some way. But chaining won't work for most basic operations in the *Layer* menu, such as Brightness-Contrast, Levels, or Curves. Operations in the *Layer* menu generally work only on the active layer.


A few operations can work either on a layer or on the whole image. *Scale* and the *Transform* submenu appear in both the *Image* menu and the *Layer* menu. Depending on which menu you use to call them, they will work on either the active layer or the whole image (all layers).

## Using Layers for Copy and Paste

Very often, you need to copy all or part of an image and put it into another image.

In the exercise at the end of Chapter 1, you dragged a small image on top of a larger one, creating a new layer. That was straightforward. But sometimes you don't have anything to drag, or you need to copy only part of an image.

For example, I have a nice image of Mars (Figure 3-23) downloaded from NASA's Hubble Space Telescope site. What would the poodle look like with Mars for a nose?



**Figure 3-23.** Mars, from NASA's Hubble Space Telescope

---

**Tip** NASA and other US government agencies, such as the National Oceanic and Atmospheric Administration, the National Park Service, and the Fish and Wildlife Service, are excellent sources for beautiful public domain images you can use in your GIMP projects. For more details, see the “Additional Resources” section of Chapter 12.

---

## Select Only the Part You Want to Paste


Pasting the whole Mars image won't work: you need to get rid of the black background first. To do that, use the Ellipse Select tool just as you did in Chapter 2.

In the red-eye project in Chapter 2, it wasn't important that the selection exactly match the pupil of the eye. But in this case, it's critical: it will be fairly obvious if some of that black background gets pasted into the poodle's face.

If you make a circular selection that doesn't end up in exactly the right place, there are two easy ways to move it. The easiest way to move a selection without moving its contents is by holding the Alt key and dragging.

On some systems, Alt-drag already has another function, such as moving the whole window. If that case, GIMP never sees your Alt-drag. If that happens to you, try holding the Shift key as well as Alt when you drag.

In addition, the Move tool has an option to move just selections rather than entire layers. The key is the *Affect*: set of buttons (Figure 3-24). The second tab, with a tooltip of *Transform selection*, makes the Move tool move only the selection, not the layer.



**Figure 3-24.** The *Affect* buttons in the *Move* tool options

If you switch *Affect tabs* in the *Move* tool to move a selection, I recommend switching back to *Transform layer* as soon as you're happy with your work. Otherwise, the next time you use the *Move* tool, it will still be in *Transform selection* mode and may give you a puzzling surprise.

If you discover after moving the selection that you also made it a bit too large or too small, you can change its size. *Select* ► *Shrink* (Figure 3-25) and *Select* ► *Grow* let you make a selection slightly larger or smaller.



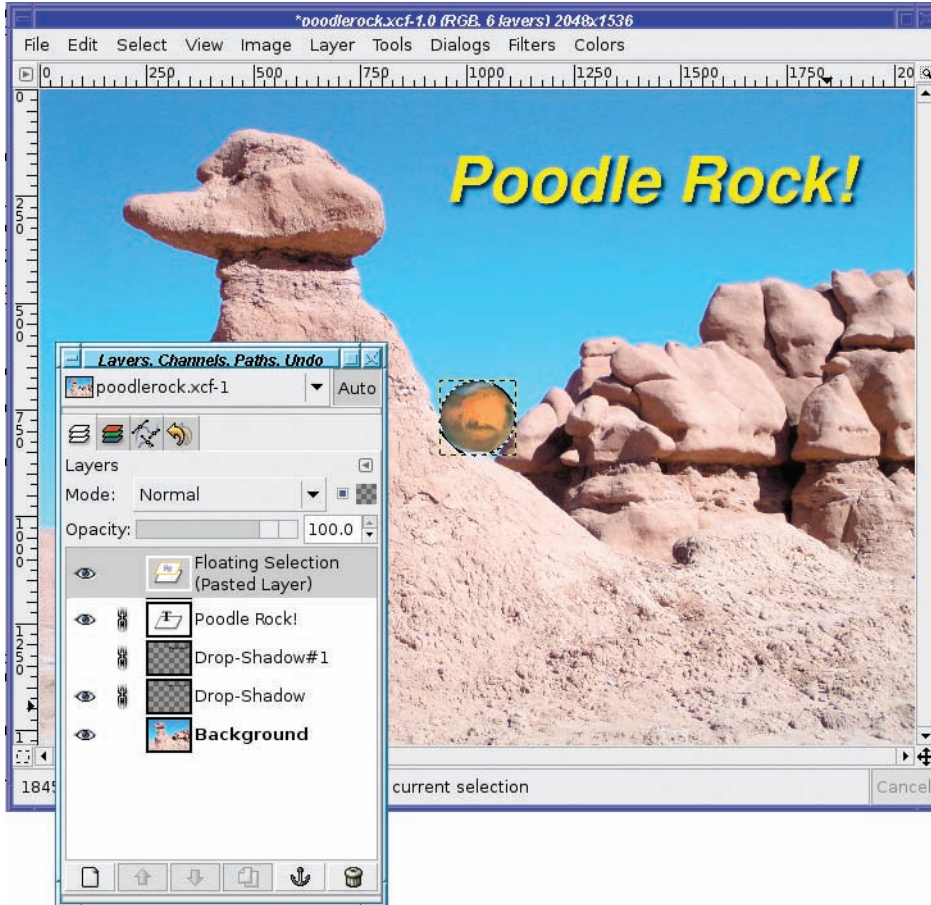
**Figure 3-25.** *Shrink Selection* dialog (*Grow* is similar)

*Shrink from image border* only matters if the selection goes all the way to the image edge. If this box is not checked, the part of a selection at the edge will not be shrunk, even though the rest of the selection will shrink.



## Paste into the Image

Now that the selection is made, you can copy in the Mars image: *Edit* ► *Copy*, or just type Control+C. Then paste into the main image: *Edit* ► *Paste*, or Control+V (Figure 3-26).



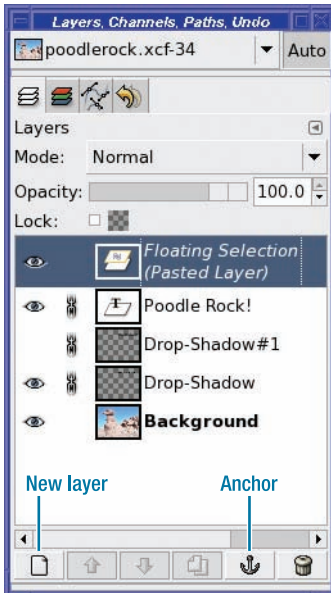
**Figure 3-26.** Immediately after pasting

A curiosity of pasting is that although the new image does paste as a separate layer, it's a special kind of layer, called a *floating selection*.

What is a floating selection?

Basically, it's a historical artifact that most people agree is not very useful any more. It may look abnormal or fragmented, many layer operations won't work right, and you can't work with any other layers while it's there.

In future versions of the GIMP, pasting may paste to a normal layer, just like drag and drop did in the project at the end of Chapter 1, rather than creating a floating selection. In the meantime, click the *New Layer* button (at the bottom of the Layers dialog, Figure 3-27) to turn the floating selection into a normal layer.



**Figure 3-27.** *New Layer and Anchor buttons in the Layers dialog*

If you're quite sure the floating selection doesn't need to be its own separate layer, and it's already in the right place, clicking the *Anchor* button will merge the newly pasted image into the previously active layer.

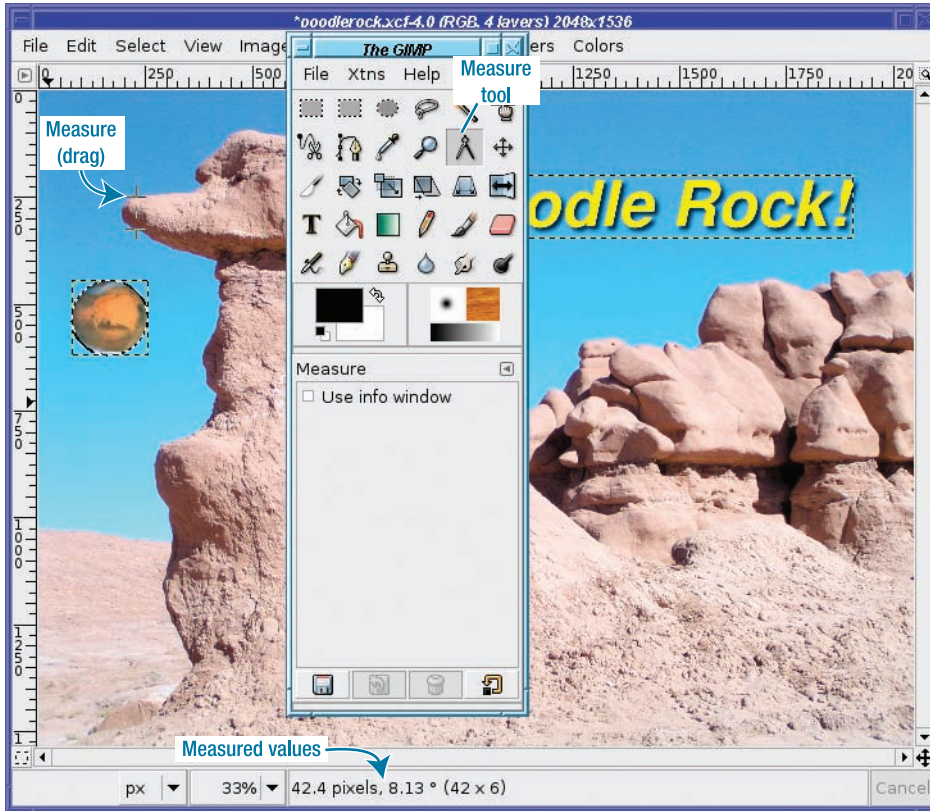
## Scale the Pasted Layer

There's one more problem—the pasted image isn't the right size. Mars is too big to be the poodle's nose!

The Scale tool in the Toolbox might be just the ticket. But if you want to make sure Mars keeps its correct aspect ratio, you might want to use the Scale dialog. You already learned in Chapter 2 how to scale an image with the dialog. Scaling a layer is just the same, except that you call *Scale* from the *Layer* menu, not the *Image* menu. But how do you figure out how big to scale it?

For that, you use the GIMP's Measure tool, an icon that looks like a set of calipers (Figure 3-28).





**Figure 3-28.** Using the Measure tool to measure the poodle's nose

Click in the image at one end of what you want to measure, and then drag to the other end. At the bottom of the image window, the Measure tool will report the distance between the two points, the angle, and (in GIMP 2.4) how big a rectangle you'd need to use to cover the diagonal.

The tool says the poodle's nose is about 75 pixels across. So I should scale the Mars layer to 75 pixels if I want it to be about the same size as the current nose. I'll make it slightly bigger, since I want the nose to stand out.

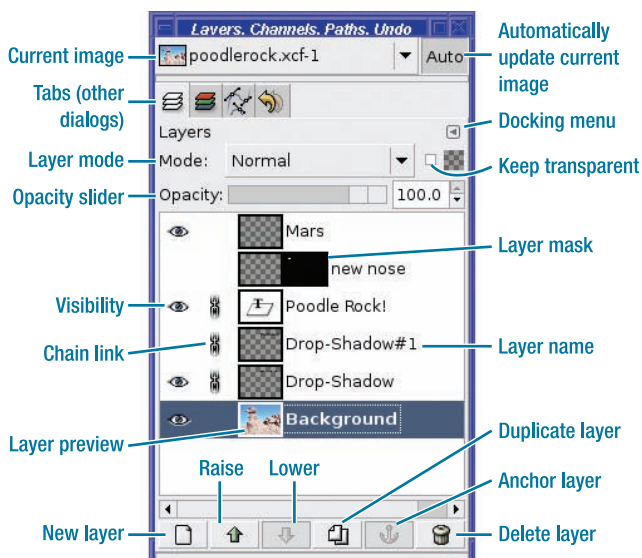
Scale the layer, use the Move tool to move it precisely where you want it—and you're done! Figure 3-29 shows the poodle with Mars for a nose.



**Figure 3-29.** *The poodle with Mars for a nose*

## A Tour of the Layers Dialog

You've seen how to do quite a few things with layers. Now it's time to explore the other features of the Layers dialog. It's small, but it packs a lot of information (Figure 3-30).



**Figure 3-30.** *The Layers dialog with all features labeled*

## Title Area

At the top of the dialog is the name of the current image, and a small preview of it. This is a drop-down menu: you can switch to a different image using this menu.

Next to it, on the right, is a toggle for making the GIMP automatically update the Layers dialog whenever you switch to another image. (You only have one Layers dialog, so it can only show the layers in one image at a time.) Usually this is what you want, but everyone has different styles of working, and some people may prefer to turn this off.

Some versions of the GIMP may not show this title area by default. To turn it on or off, click on the icon for the Docking menu (see next section), and choose *Show Image Selection*.

## Tabs

Below the current image menu are the tabs for any other dialogs that may be docked in this window (remember “Docking” from Chapter 1?). Typically Channels, Paths, and Undo History are docked in the same window with Layers, though Layers is by far the most commonly used.

If you don’t have any other dialogs docked in the window, the tab area won’t be visible.

The *Docking menu*, on the right just below the tabs, lets you close or un-dock the dialog, or configure the dialog window in other ways. The word “Layers” to the far left of the Docking menu is simply the name of the current tab.

Some versions may also show an *x* button next to the Docking menu, a shortcut to remove the dialog from the window. If other dialogs are docked in the window, the window itself won’t close—just the current tab.

The rest of the dialog applies to whichever layer is currently active.

## Layer Mode

First is *Layer mode*.

So far you’ve only seen layers combined in Normal mode: anything that’s opaque in the top layer is visible, while anything transparent lets you see through to lower layers. However, layers can be combined in many other ways, such as Overlay, Multiply, Dodge, or Burn, in which you see the two layers combined according to various mathematical functions. See Chapters 9 and 10 for more information on how to use other modes.

Next is a button called *Keep transparent*. (The checkerboard icon next to it is similar to the checkerboard background the GIMP uses to indicate transparency, and is simply a reminder of the button’s meaning.) Its position varies with the GIMP version: it may be to the right of the *Layer mode* menu or below it, but look for the checkerboard icon.

When *Keep transparent* is checked, you will not be able to draw anywhere that the current layer is transparent. This is very useful for drawing, but is also a common source of confusion. You’ll learn more about this option in Chapter 4.

## Opacity

Next is the *Opacity slider*.

In addition to having transparent areas, a layer can be entirely transparent, entirely opaque, or anything in between. You’ve seen that indirectly when making drop shadows: the transparency entry in the Drop-Shadow dialog sets the opacity for the drop shadow layer.

Select your drop shadow layer in the Layers dialog, look at the opacity slider, and notice that it's at something less than 100%. Try changing the opacity, and watch the effect it has on the appearance of the shadow.

You can use the opacity slider to create all sorts of useful effects. For example, you can render a text layer translucent and make it float over an image. Or you can create a layer that's entirely white, make it mostly transparent, and lay it over the top of an image to create an impression of haze. That's just the tip of the iceberg.

## Layers List

Below the opacity slider is the list of layers in the image (with a scrollbar if needed), one layer per line. This is also called the *layer stack*, because the layers are stacked one above the other. You're already familiar with most of the items shown, but there are a few you haven't seen yet.

### The Visibility Eye

On the far left is the “eyeball” *Visibility* icon. When it's on, the layer is visible; when it's off, invisible. Clicking on the eye toggles the layer's visibility. Shift-clicking makes *only* that layer visible; another Shift-click brings all the other layers back.

### The Chain Link

Next is the *Chain link* icon. When several layers display this icon, they are linked together. Moving or transforming any one of the layers (via tools such as Rotate) will work on all of them together.

Shift-clicking on the chain link removes any layers that were already linked. This can offer a shortcut if you have a lot of layers linked and want to unlink them.

### The Layer Preview

The *Layer preview* shows a small thumbnail image of what the layer looks like.

You can't tell the difference between two text layers using the preview, but it can be useful for telling image layers apart. Double-clicking on the preview brings up a dialog that lets you edit the layer attributes (mostly, its name). Clicking on the layer preview is one way to select the layer (make it active). It also ensures that if you draw, you will draw on the layer and not its mask.

Its mask?

### Layer Masks

To the right of the layer preview is the *Layer mask preview* (if there is a layer mask).

A mask controls which part of the layer is visible. Most layers don't have masks, but when a mask does exist, wherever it is white, the layer is visible. Where the mask is black, the layer will be invisible (transparent) even if that part of the layer has something in it.

Clicking on the layer mask makes it active: if you draw with the mask selected, you'll be drawing in the mask, not in the layer. On a selected layer, the border of either the layer preview or the layer mask preview will be highlighted to indicate which one is active.

A layer mask is another way of representing transparency. In some cases it's easier to work with a mask than to work with transparency directly. You'll learn how to use the basics of layer masks in Chapter 5.

## Layer Name

To the right of the layer and mask previews is the name of the layer. Clicking on the layer name selects the layer; double-clicking allows you to edit the layer name (for example, to change “Pasted Layer” to “Mars”). You can change layer names at any time.

## Layer Buttons

Finally, at the bottom of the Layers dialog is a row of buttons that can create or delete layers, or change the order of existing ones. You can either click on them (to affect the active layer) or drag a layer on top of them. They all have tooltips, so you can remind yourself which button does what. The buttons each correspond to an action in the image window’s *Layer* menu; some of them also have key bindings, which you can learn, or set, using the *Layer* menu.

### The New Layer Button

*New layer*, the leftmost button, creates a (surprise!) new layer. You can choose the layer name, size, and fill color. (You’ll create lots of new layers in Chapter 4 when you start making drawings.)

Dragging an existing layer onto the *New layer* button creates a blank new layer with the same size and position.

New layer can also convert a floating selection into a regular layer.

### The Raise and Lower Buttons

The *Raise* and *Lower* buttons move the current layer up or down in the layer stack. If you want something to appear *on top of* something else, adjust the orders of the layers with these buttons.

You can also change the order of layers by dragging a layer preview to a new place in the layer stack. Not only that, you can drag a layer preview out of the Layers dialog into a new image, to add a copy as a new layer in that image. You can even drag to the Toolbox, to create a new image containing only that layer.

### The Duplicate Layer Button

*Duplicate layer* creates a new layer that is an exact copy of the active layer, or of any layer you drag onto the button. You can then move the layer up or down in the layer stack, or modify it in all kinds of ways. You’ll see some ways to use this when you create an animation, at the end of this chapter.

### The Anchor Layer Button

*Anchor layer* is used to merge a floating selection with whichever layer was previously active.

### The Delete Layer Button

*Delete layer* deletes the active layer. You can also drag a layer to the button to delete it.

---

**Tip** If you change your mind about a layer operation and need to undo, if Control+Z in the Layers dialog doesn’t undo it, try undoing in the image window.

---

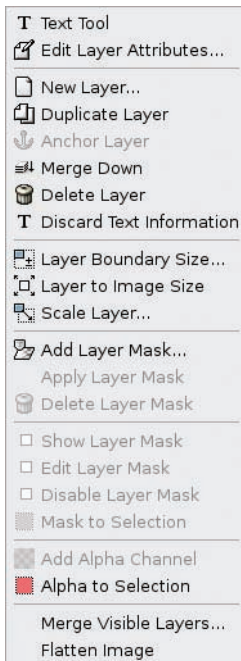
## Layer Context Menus

Right-clicking on any layer line brings up a context menu offering operations on that layer (Figure 3-31).

---

**Caution** Right-clicking on a layer line will also make it active. If you're not careful, it's possible to end up working on the wrong layer.

---



**Figure 3-31.** *The layer context menu*

### The (Sometimes) Text Tool Item

If the layer is a text layer, the top item will be *Text Tool*. Choosing this will display the Text Editor window, so you can make changes to the text. It also activates the Text tool options in case you need to make changes to color, font, size, or other text properties.

If the layer is not a text layer, this item won't be there.

A text layer is a special type of layer, indicated by a special icon in place of the normal layer preview. Many changes to text layers, such as drawing on them or rotating them, change the layer to a normal graphics layer. After that, you will no longer be able to edit the text in the Text Editor window, nor change the text properties.



## Edit Layer Attributes

*Edit Layer Attributes...* brings up a dialog showing some properties of the layer you can change. Usually all this offers is the layer name. It's easier to change the layer name by double-clicking on the layer's name in the layer stack, so this item isn't very useful. Double-clicking on the *Layer preview* button is another way to bring up this dialog.

## Layer Controls

*New Layer...*, *Duplicate Layer*, *Anchor Layer*, and *Delete Layer* have the same meanings as the corresponding buttons.

*Merge Down* combines the active layer with the layer below, using the active layer's mode. Any transparency is retained.

*Discard Text Information* will only appear for a text layer. It changes a text layer to a normal graphic layer. Lots of other operations will also do this as a side-effect. However, there's not much reason to do it before you need to.

## Layer Boundary Size

Every layer has a size, which may be smaller than the whole image. For example, text layers are just barely big enough to hold the text they contain. *Layer Boundary Size...* lets you change a layer's size.

Most often you'll use this to make a layer slightly larger—for example, to make room if you want to blur a text layer, make a line longer, or increase the size of a white background area.

## Layer to Image Size

*Layer to Image Size* is a special case of *Layer Boundary Size*: it makes the layer as large as the whole image. You often need to do this after increasing image size.

## Scale Layer

You've already seen *Scale Layer...*, which makes a layer and its contents larger or smaller.

## Mask Options

*Add*, *Apply*, *Delete*, *Show*, *Edit*, *Disable Layer Mask*, and *Mask to Selection*, are all tools for working with layer masks which you'll learn about in Chapters 5 and 9.

## Add Alpha Channel

*Alpha* is a fancy term graphics people use to mean transparency. *Add Alpha Channel* makes the layer capable of using transparency. It doesn't make anything transparent immediately, but if you use *Edit* ► *Clear*, or use tools such as the Eraser, you will create transparent areas instead of white.

If you see an error when you try to raise the bottom layer in the stack, or the arrow to raise the layer is grayed out, it may be because it needs an alpha channel.



## Alpha to Selection

*Alpha to Selection* selects everything in the current layer that is *not* transparent. The name is a bit confusing, since it suggests that only the transparent parts would be selected.

## Merge or Flatten

*Merge Visible Layers...* and *Flatten Image* are two different ways of combining multiple layers into one.

*Merge* retains any invisible layers, and shows a dialog asking what to do about the layer size of the result.

*Flatten* merges all the visible layers, deletes any invisible layers, fills any transparent areas with the current background color, and gives you an image with a single layer and no alpha channel. *Flatten* is sometimes done for you, temporarily, when you save to a format such as JPEG that cannot represent transparency or layers.

## Bonus Project: Making Simple GIF Animations

Animations are a fun use of layers.

An animation is a set of images in which the picture changes slightly from frame to frame. Played one after the other, they make a movie.

The GIMP represents each animation frame as a layer. The first frame is the lowest layer; the last frame is the top layer.

Since you know how to create layers, there's not much more you need to know to make an animation. So let's jump right in and try one.

I have a nice photo of a biplane (Figure 3-32). Wouldn't it be a kick to make it land and take off?



**Figure 3-32.** A biplane. You can make it fly!

## Create a Base Image

First, you need to create an image big enough for every part of the animation. About ten times the length of the plane, and four times the height should give you plenty of room. (You can always crop it later if you don't need that much room.)


Usually you'll want to scale animations fairly small, especially if they'll be put on the web. An animation stores quite a few images in one file, so the file size can be quite large even if the pixel size isn't. If you're animating an image that starts as a high resolution digital photograph, you'll probably want to scale it down quite a bit. Also, a flat (single-color) background will compress to a smaller file size.

But first, before you create the new image, why not give it a background the color of the sky?

*New Image* (unless you tell it otherwise) will use the GIMP's current background color as the color of the new image. So click on the background color swatch in the Toolbox, and choose a nice light blue similar to the background of the plane.

To match a color in an existing image, such as this plane's blue background, you can use a GIMP tool called *Pick Colors from the Image* (Figure 3-33). By default it sets the foreground color. In this case, we want it to set the background color, so check that box in the tool options.

It also pops up a dialog with information about the color (which you can read, ignore, or dismiss).



**Figure 3-33.** Picking colors from the image. Notice that Set background color has been checked.

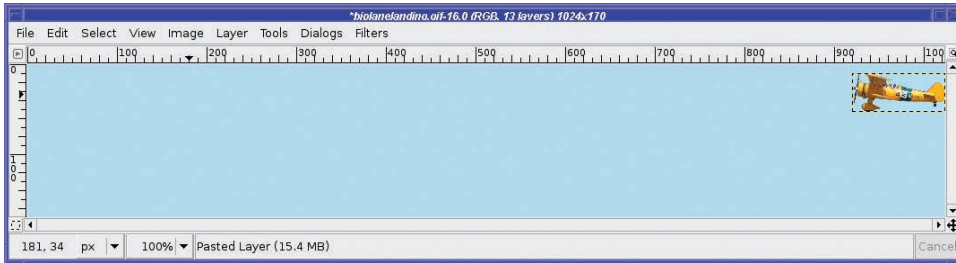
Once the background color has been set, it's time to create the new image, with *File* ► *New...*, at whatever size is best for your animation.

## Add the Animation Frames

Now comes the time-consuming part: adding each frame of the image.

Start by copying (*Edit* ► *Copy* or Control+C) the biplane image. Paste into the new image, click the *New Layer* button to turn the floating selection into a regular layer, and then use the *Move* tool to move it to where you want it for the first frame (Figure 3-34).

Remember, the order of layers in the stack is the order in which they'll be played in the final animation, starting with the lowest layer.

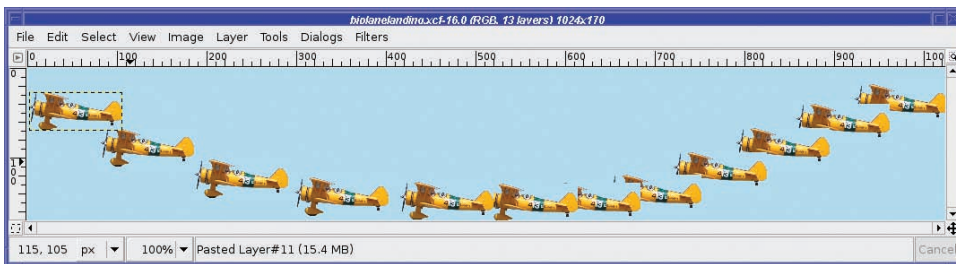


**Figure 3-34.** *The first frame has been pasted.*

If you've gotten into good layer management habits, you may be tempted to name each layer, e.g., "biplane 1", "biplane 2", and so on. Although that's usually a very good idea, it isn't worthwhile at this stage. The layer names will change, perhaps several times, before we're through.

Next, paste the same image again. (You shouldn't need to copy again first.) You may think that the second paste didn't work—there's no visual indication that anything happened. Why not? Because GIMP pasted the second image right on top of the first. You can tell by looking at the Layers dialog that you have a new floating selection in addition to the biplane layer you pasted already. If you drag the new layer (using the Move tool) to its new location, you'll see there's another identical biplane layer right underneath it.

Continue pasting, clicking *New Layer*, and then moving the pasted layer to a new position, until you've made as many animation frames as you want. Take a look at your creation and decide if it looks like you want it. Your image may look something like Figure 3-35.



**Figure 3-35.** *All the planes pasted*

---

**Note** You can also use *Duplicate Layer* to make each new layer, rather than successive paste operations.

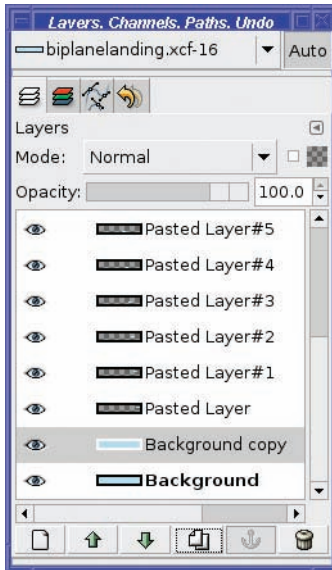
---

Save your work (as an XCF) now. There's potential for error in the next step, so make sure you have a copy of what you've already done.

## Give Each Frame Its Own Background


The final step is to add in the sky on each frame. The frames depict a biplane on a transparent background. But in the final animation, each frame will replace the previous one, so each frame needs its own copy of the sky.

This is tedious, but straightforward. In the Layers dialog, select the background layer and do a *Duplicate layer*. This will create a new layer called “Background copy” (Figure 3-36).



**Figure 3-36.** Create a copy of the background.

Now right-click on the first layer above “Background copy” and select *Merge Down* (Figure 3-37).

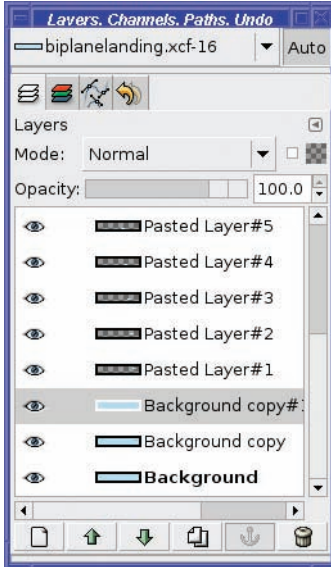


**Figure 3-37.** Merge the first pasted layer with the Background copy.

At this point, the first biplane layer has a background of its own, and the combined layer is called “Background copy.”

Now go to Background and make another duplicate (“Background copy #1”). This time, use the *Raise* button once to raise it so that it’s underneath the next layer which still has transparency (Figure 3-38).

You can tell which layers have transparency from looking at the layer thumbnail. The checkerboard pattern GIMP uses to indicate transparency is visible in the thumbnails, as in Figure 3-38.



**Figure 3-38.** Layers dialog after merging the first layer into “Background copy” and raising “Background copy #1” to be under “Pasted Layer #1”

Then right-click on the layer with transparency (probably “Pasted Layer #1”) and *Merge Down*.

Continue this process, making a copy of the background layer, moving it up to just below the first layer that still has transparency, and then merging the transparent layer down until you get to the last frame.

If you find it easier, of course you could create the appropriate number of background copies all at once, move them to the right spots, then merge each plane down into its respective background. It’s a matter of preference.

You may notice that at each step, one more plane disappears in the image the GIMP is showing you. This is because each time you add the sky background, that layer becomes opaque, so you can no longer see the airplanes in the layers below it. Clicking the visibility button on a few layers should convince you that your planes are still there.

Once you’re finished merging the sky into all your frames, you should probably turn off visibility on the background frame or delete it, unless you want your animation to start with an empty blue sky.

## Test Your Animation

Check your work: the GIMP should show only the last frame, but if you turn visibility off for the top layer, you should see the next frame down; turn off visibility for that layer too, and you'll see the next layer; and so on.

---

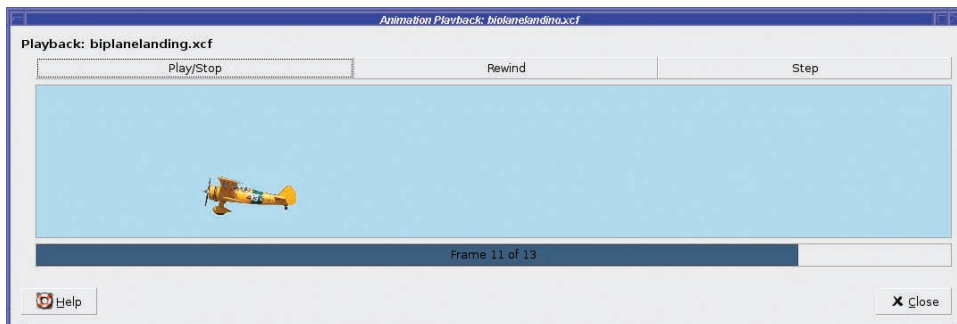
**Tip** You can also Shift-click on a layer in the Layers dialog to see only that layer, and turn off visibility on all other layers. Another Shift-click brings the other layers back.

---

Now turn them all back on. It's time to test the animation!

The GIMP has a very basic animation player in *Filters* ► *Animation* ► *Playback...* (Figure 3-39). It has a play/stop button, rewind button, and step button. But that should at least give you an idea of how well your animation is working.

Don't worry if it's too fast. You can control that in the next step.

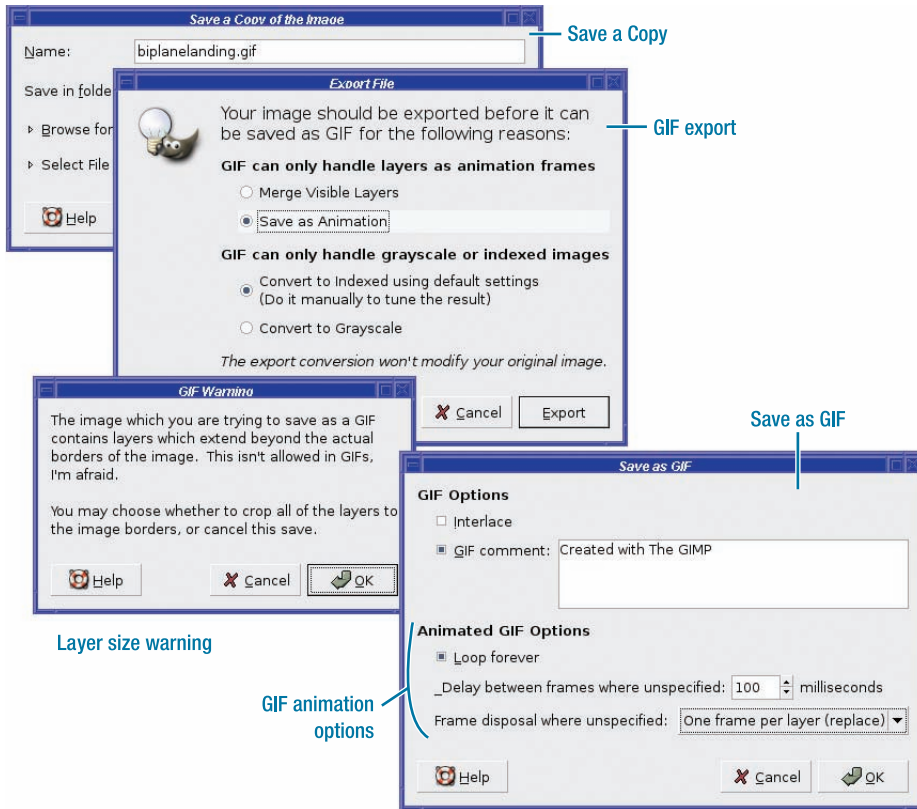


**Figure 3-39.** Animation playback

## Animation for the Web: Save as GIF

Now, if you save a copy of the image as a GIF, you get a series of dialogs (Figure 3-40).





**Figure 3-40.** The dialogs involved with saving a GIF animation

First is the normal Save a Copy dialog; type a file name ending in “.gif” here, or use the menu to select the GIF file type.

This leads to the GIF Export dialog. It reminds you there are multiple layers, and asks you whether to merge all the layers into a single still image or create an animation. Merging is the default; you will have to check the box for animation each time you save, unfortunately.

The GIF Export dialog is also the one controlling conversion to indexed mode. You can make your animation file smaller if you convert the image to indexed mode first, as was discussed in Chapter 2.

If you’ve dragged any of your layers beyond the image borders, you may get a warning like the one shown. It doesn’t do any harm; if you see this, just click OK.

Finally you get the important dialog: Save as GIF. This offers several animation options.

*Loop forever* is on by default. You can make your animation play once then stop, or loop back to the beginning once it gets to the end.

*Delay between frames* controls the speed of the animation. The default is 100 milliseconds (1/10 of a second), but you can make the animation run slower or faster if you choose.

*Frame disposal* can usually be set to *I don’t care*, as long as you’ve made each frame opaque, as in this example. If you *do* care, the choices are *Cumulative layers (combine)* and *One frame per layer (replace)*. This option only makes a difference if you make an animation with transparent frames: it lets you control whether previous frames are erased before drawing the next.

## TUNING YOUR GIF ANIMATION WITH LAYER NAMES

Remember I suggested not changing the layer names early on? You saw how the layer names got lost when you merged them with copies of the background. But once you've done that, it's safe to change the layer names to something a little clearer.

While you're looking at layer names, though, it's useful to know that layer names can also specify the interval between frames. You can even vary the intervals: have your biplane come in slowly, and then use a smaller interval as it's revving up to take off.

To specify frame rate of a layer, edit the layer name (by double-clicking on the layer name in the Layers dialog) and give the layer a name that ends with a millisecond interval (in parentheses). The name might look something like: "biplane 3 (200ms)."

You can also add (*combine*) or (*replace*) after the layer name to indicate that the current layer should combine with what's in the previous frame without replacing it, or should replace everything. The name now might look like: "biplane 3 (200ms)(replace)."

If you use *File* ► *Open* to open the animated GIF you saved, and look at the Layers dialog, you'll find that the GIMP will now show the frames with names such as "Frame 8 (200ms)." Rather than edit each frame to insert the interval, it's usually easiest simply to open the GIF, and then change the few layers where you want to use a different interval.

*Filters* ► *Animation* ► *Optimize (for GIF)* will also insert intervals in each layer name; but it's hardwired to 100ms, so it's not very helpful if you want most frames to have a different interval.

Now that you know how to make a biplane land and take off again, it's not *that* much harder to make it loop. You'd need to start with a taller blank image, of course (or use *Image* ► *Canvas Size*, which you'll learn about in Chapter 4). Then use the Rotate tool on each layer before you move it to the right place. I'm sure you can think of other fun animation projects as well. The sky's the limit.

## Summary

By now you should be very comfortable with using layers to make new images.

You keep your Layers dialog visible and you know how to use it. You've used layers for newly created text, and for objects pasted or dragged from other images. You know how to link several layers together to move them as a group, and apply transforms or other GIMP operations to a single layer. In addition, you've learned some basic GIMP tools such as the Move tool and the color chooser.

You're ready for anything. With that in mind, it's time to lighten up a bit, take what you've learned about layers, and apply that knowledge to a really fun aspect of the GIMP: drawing.





# Drawing

In this chapter, I'll take a step back from photographs and explore the GIMP as a drawing program. I'll cover the following topics:

- Drawing lines and curves
- Changing colors and brushes
- Erasing
- Drawing rectangles, circles, and other shapes
- Outlining and filling regions
- Filling with patterns and gradients
- Importing brushes or gradients, or making your own
- When GIMP won't draw
- A drawing project

## A New Image

Most drawing projects begin with a blank canvas in a new window. Create one using *File* ► *New* from either the Toolbox or any image window.

Choose a reasonable size for the image. The default values in the New Image dialog will reflect the current image or selection. To practice drawing techniques, something that fits on your screen is best: I often use 800×600. For real-world projects, you may want to use a larger canvas, in order to have high resolution for printing or for including fine details. The *Template* drop-down menu at the top of the New Image dialog lets you choose from a list of popular image sizes.

---

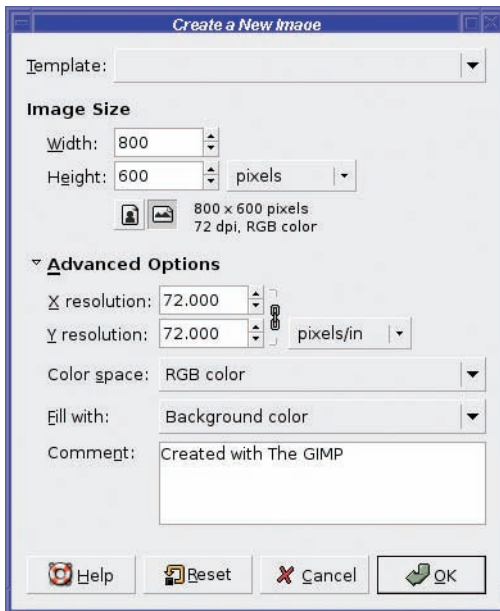
**Caution** In the Preferences category *Default Image* (called *New Image* in earlier versions of the GIMP), you can set an image size. However, the New Image dialog may not always show this size. It may reflect the dimensions of the last image you created, or of the last region you copied. To go back to the default size you specified in Preferences, click *Reset*.

---

Expanding the *Advanced Options* tab in the New Image dialog (Figure 4-1) shows some additional choices, such as

- *Resolution*
- *Color space* (RGB versus grayscale)
- The color to make the new image (*Fill with*)

By default, GIMP will fill your new image with the current background color. For starters, you probably want a white background; though later you might want to use other colors for a richer effect. You can change your background color using the color swatch in the Toolbox, or for white, you can simply choose *White* from the *Fill with* menu of the New Layer dialog.



**Figure 4-1.** *The New Image dialog showing Advanced Options*

You can also make a new image transparent. This can be useful if you're making an icon for a web page, or an image that will be pasted onto a photograph later. But the gray checkerboard pattern GIMP uses for transparency can be distracting. When you're learning drawing techniques, it's easier to draw on layers over a white background, and then turn off the background before you save the file. You'll see how that helps in the drawing project at the end of this chapter.

The New Image dialog offers one more option: a space for a comment, defaulting to *Created with the GIMP*. You can use this space for your name and copyright information, or information about the image you're creating and what it represents. Or, of course, you can just leave it blank.

---

**Note** Not all image formats can include a comment, but the most common formats—JPEG, PNG, and GIF—all have it.

---

## Using Layers for Drawing

The first rule of drawing is: Use a new layer.

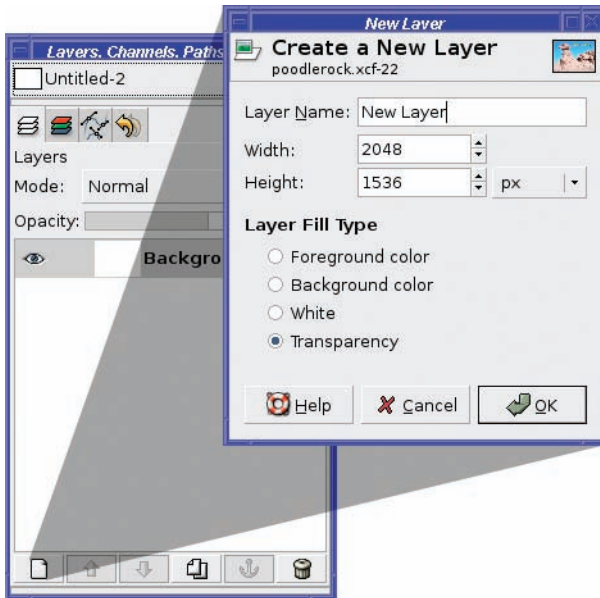
You just created a new image with a perfectly clean background. Why should you add yet another layer?

What if you want to change the background color later? Or make the background transparent? What if you decide you want to move part of your image to a different place? What if you want to duplicate a figure you've drawn, so you can have two of them against the same background?

Sure, it's possible to make all these changes with the GIMP later: but it's a lot more work than just using layers to begin with.

You don't need to make a new layer for every line you draw. But try to think of your drawing in terms of functional units: the background is one layer, grass might be another, trees a third, and the sky a fourth.

So, the first step is to create a new blank layer. To do that, go to the Layers dialog and press the *New Layer* button in the lower-left corner. Up pops the New Layer dialog (Figure 4-2).



**Figure 4-2.** *The New Layer dialog*

The *Layer Name* field lets you choose a memorable name for the layer. You don't have to choose a name; the GIMP will assign one for you, something like "New Layer." But a well-chosen layer name can help in the long run. The name shows up in the Layers dialog, so a layer named "red arrow" or "sky" can tell you right away what's in it.

*Width* and *Height* default to the size of the image. This is usually fine. If you know you don't need a layer quite that big, you might be able to save some memory by specifying a smaller size.

*Layer Fill Type* lets you specify whether the layer will start out transparent or not (much like *Fill with* in the New Image dialog). The default is *Transparency*. Most often, when you add new drawing layers, transparency is just the ticket. But some of the special effects you'll learn in later chapters will use a solid-colored layer.

The *OK* button creates the new layer and adds it to the Layers dialog. If you need to move it up or down in the layer stack, use the up and down arrow buttons at the bottom of the Layers dialog, next to the *New Layer* button; or drag the layer's line to where you want it to be. The new layer automatically becomes the active layer.

---

**Note** If you tried moving your new layer down in the stack (the only choice available now) you were probably surprised to discover that you can't! Remember the note about right-clicking to *Add Alpha Channel* in the Layer's dialog context menu in Chapter 3? By default, the background layer doesn't have transparency and can't be moved up.

---



## Drawing Lines and Curves

You're ready to draw! Choose a nice color (using the foreground color swatch). Or just leave it black.

The GIMP has a collection of four tools for drawing lines and freehand curves. I'll start with the simplest: the Pencil.

### Drawing Hard Edged Lines: The Pencil Tool

The Pencil tool lets you draw sharp-edged lines (Figure 4-3). Select it in the Toolbox, and then try scribbling by dragging on the image to see what it does. The Pencil tool will leave a trail everywhere you go, as long as you have the left mouse button pressed.

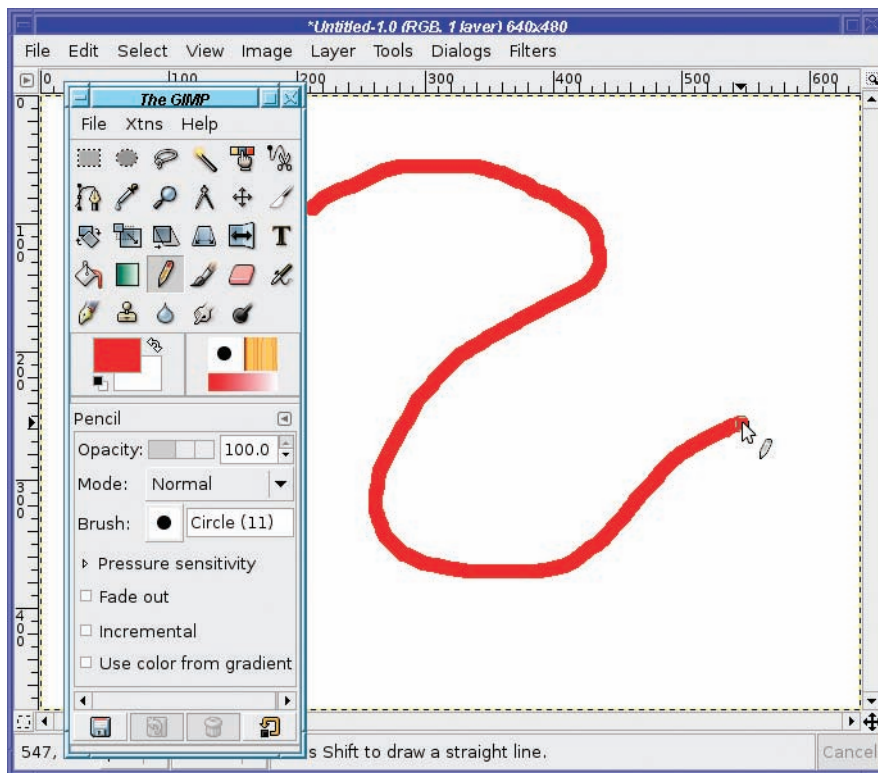


Figure 4-3. Using the Pencil tool

---

**Tip** It's easy to draw straight lines with any of the line-drawing tools. Click once where you want one end of the line to be (GIMP draws a circle there). Then move (with the mouse button up—you're not dragging) to where you want the other end of the line, hold down the Shift key (GIMP will show a thin line between the mouse position back to the last place that you clicked), choose your end point, and when you're happy with the line, click the mouse button.

---

Notice the mouse cursor in Figure 4-3. In addition to the normal arrow and the pencil icon telling you which drawing tool is selected, there's a circle around the arrow's point. This shows the size of the current *brush*.

Okay, a pencil normally has a point, not a brush. But this is a Super GIMP Pencil (see the sidebar “GIMP Brush Icon Display Options” for more information).

### GIMP BRUSH ICON DISPLAY OPTIONS

You can stop GIMP from showing brush size in the cursor by turning off *Show brush outline* under *Preferences* ► *Image Windows*.

But this feature is very useful: most people, except on extremely slow machines, will want this enabled. *Show paint tool cursor*, in the same Preferences screen, will turn off the pencil and arrow icons and show *only* the brush outline.

Alternately you can turn off *Show brush outline*, turn on *Show paint tool cursor*, and change the cursor mode: from *Tool Icon* to either *Tool icon with crosshair* or *Crosshair only*. This gives you a similar but perhaps more precise result. Your call! You can also choose *Black and White* instead of *Fancy* for *Cursor Rendering*, which may help performance slightly on very slow machines.

## Brushes

Changing the brush gives you control over the width and shape of the line you draw with any of GIMP's drawing tools. There are two ways of changing the brush. You can click on the brush icon in Tool Options, which drops down a menu (Figure 4-4). Or you can click on the brush icon in the Toolbox, to the right of the color swatches, which pops up the Brushes dialog (Figure 4-5).

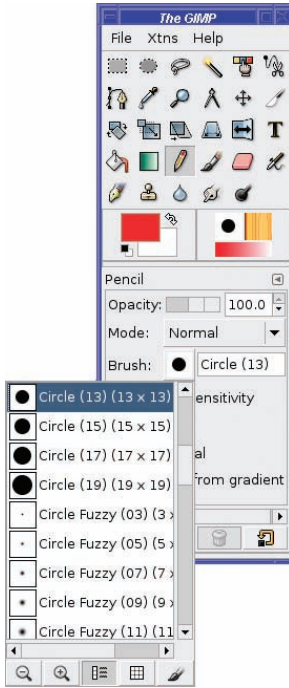


Figure 4-4. Choosing a brush from the drop-down menu

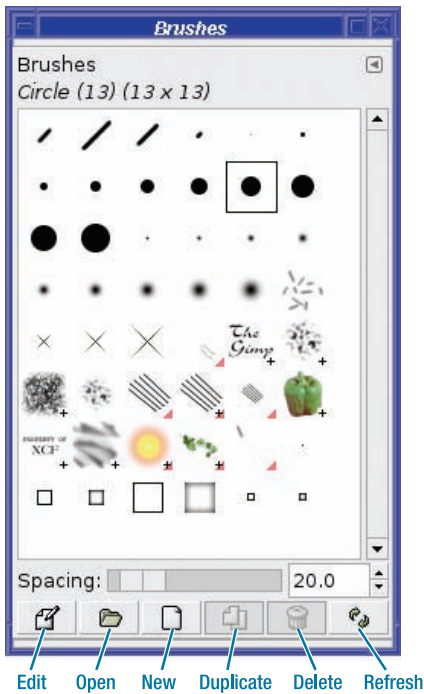


Figure 4-5. The Brushes dialog

Clicking on a new brush in either the menu or the dialog changes the active brush immediately. You can try one, draw with it, undo it if you don't like it, and try another brush, all without dismissing the dialog.

---

**Tip** You can also get to the Brushes dialog using *File* ► *Dialogs* ► *Brushes* from the Toolbox, *Dialogs* ► *Brushes* from an image window, or from various docking menus, since it's a dockable dialog.

---

Brushes are like very small images. When you use a paint tool such as the Pencil, it's as though you dipped the brush image in ink and then dragged it across the screen.

In addition to the normal hard-edged circles, some brushes have fuzzy edges. The Pencil tool will ignore the fuzzy edge; I'll talk about how to use those brushes in a moment.

Some brushes, such as the diagonal slashes, are asymmetric. They yield different patterns depending on the direction in which you drag. (You can use this effect for calligraphy, as if you were using a classic quill pen.)

A few brushes are actually animated images. Why would you ever want an animated brush? Since it changes as you drag it across the screen, you get a line that's varied or random. A fun example is the Vine brush (Figure 4-6).

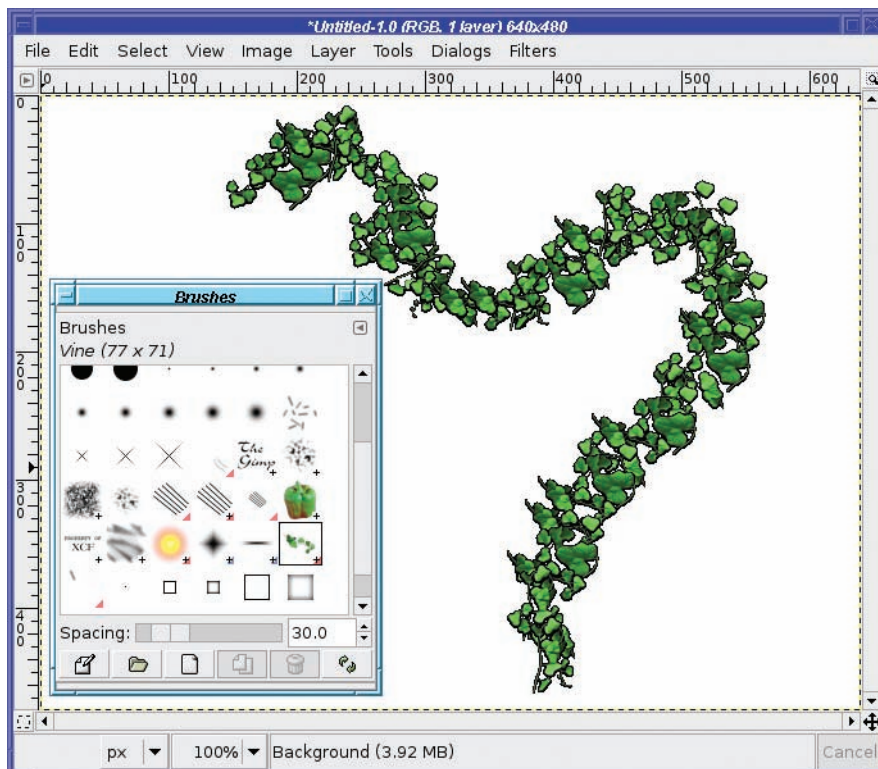


Figure 4-6. The Vine brush

---

**Note** Like a few other colored brushes, the Vine brush has another weird property: a fixed color that ignores GIMP's current foreground and background colors.

---

The Brushes dialog has a few options, shown in Figure 4-5. *Spacing* can make the brush's patterns look more spread out as you drag across the screen. Try it with the Vine brush to see the effect, but you can use it on non-animated brushes as well.

Most of the buttons at the bottom of the Brushes dialog help with changing brushes or adding new ones.

- *Edit brush* brings up a Brush Editor dialog. On most existing brushes, everything in this dialog will be grayed-out because built-in brushes cannot be edited: they're read-only.
- *Open brush as image* (new for GIMP 2.4) does, well, the obvious. You can edit the image, and then save it as type *GIMP brush* (with the extension “.gbr”) or *GIMP brush (animated)* (“.gih”) in the Brushes folder of your GIMP profile. In the case of animated brushes, you'll be able to see each frame as a layer, just like the animation you made in Chapter 3. You'll see a more detailed example of making an animated brush in Chapter 9.

---

**Tip** Because a brush is just an image, you can create a brush image yourself in GIMP, starting from scratch. You don't have to start by editing an existing brush.

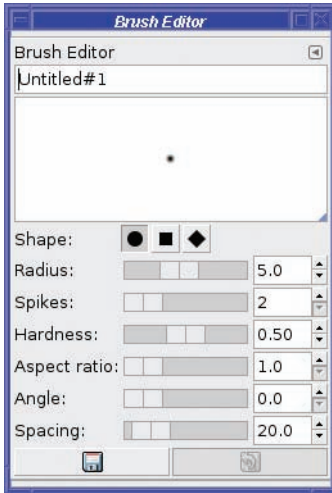
---

- *New brush* opens the Brush Editor dialog and lets you create a *parametric brush* (Figure 4-7). You can specify shape, size, hardness, and related parameters. The button at the lower-left lets you save your brush by the name chosen at the top of the Brush Editor. Your new brush will then appear in the Brushes dialog every time you start the GIMP, unless you delete it.

---

**Note** When you save a parametric brush, the file will appear with a .vbr extension in your GIMP brushes folder.

---



**Figure 4-7.** Creating a parametric brush with the Brush Editor

- *Duplicate brush* makes a copy of a parametric brush which you can change, and then save as a different brush.
- *Delete brush* deletes a brush. You can't delete any of the built-in brushes, only those that you saved in your personal collection.
- *Refresh brushes* re-reads your brushes directory, to see if you've installed any new brushes while GIMP was running.

---

**Tip** There are lots of excellent GIMP brushes available for free downloading. Try a web search for *gimp brushes*.

---

## Parametric Brushes vs. Image Brushes

Most brushes are images. But a brush you create with the Brushes dialog's *Edit* or *New* buttons is a different type, as mentioned previously, called a parametric brush. You can choose from a few simple shapes (round, square, diamond), and then specify the size (*Radius*), *Hardness* (how fuzzy the edge is), *Aspect ratio* (whether it's long and thin or basically square), *Angle* (which rotates it), and *Spacing* (which makes it draw discrete blobs as you drag across the window, instead of placing the images close together so that they appear to be a continuous line).

Since parametric brushes and image brushes are specified differently, you can't change one into the other. You can't click *Edit brush* on an image brush, and you can't click *Open brush as image* on a parametric brush. You don't need to remember which is which, since GIMP will gray out the buttons that don't apply to the current brush.

In the Brushes dialog, each parametric brush shows a blue triangle at its lower-right corner, while animated brushes have a red corner.

## Drawing Tool Options


All the line-drawing tools have similar options. The basics shared by all tools include the following:

- *Opacity* makes the tool's line more transparent. It's a percentage, from 0 to 100.
- *Mode* offers the same list of modes as the Layers dialog (which will be discussed in Chapters 9 and 10), plus a few that are special (see below, "Special Drawing Tool Modes").
- *Fade out* makes the line fade out after a specified distance, even if you keep dragging.
- *Incremental* only makes a difference if *Opacity* is less than 100%. Using a tool with no spacing, or drawing on top of a previous line, will make the line darker (more opaque). Without incremental mode, the line won't get darker.
- *Use color from gradient* uses the gradient (displayed below the current brush in the Toolbox window) instead of the current foreground color. Click on the gradient in the Toolbox to see some examples; you'll learn more about gradients later.

In addition to these options, the drawing tools can respond to pressure and tilt if you have a drawing tablet that can send these signals.

### Special Drawing Tool Modes

In addition to the normal layer modes (which will be discussed in Chapters 9 and 10), the drawing tools offer three modes that are special (Figure 4-8).



**Figure 4-8.** *The three special drawing modes*



- *Dissolve* adds randomness to the drawing. Wherever GIMP would have drawn partial transparency, instead it will draw a random pattern of dots. (*Dissolve* is also available as a layer mode, but it's more often useful for drawing.)
- *Behind* draws behind anything that's already drawn in the layer. This only works when drawing onto a layer that has transparency. (On an opaque layer, such as a background layer filled with white, it would make no sense. You wouldn't be able to see something drawn *behind* an opaque background!)
- *Color erase* seeks out the current foreground color and erases it, replacing it with transparency. It doesn't erase any other colors; only the current foreground color. Of course, like *Behind*, this only works if the layer allows transparency. In Figure 4-8, the foreground color was changed to red before drawing the color-erase sample; it would have had no visible effect when the foreground color is black, since it doesn't match the red box.

## Drawing Fuzzy or Smooth Lines: The Paintbrush Tool

The Pencil tool is okay when you're using hard-edged brushes. But what if you want to use one of those fuzzy brushes? Or what if you want a line without those “jaggy” edges that the Pencil tool makes?

In that case, use the Paintbrush tool (Figure 4-9).

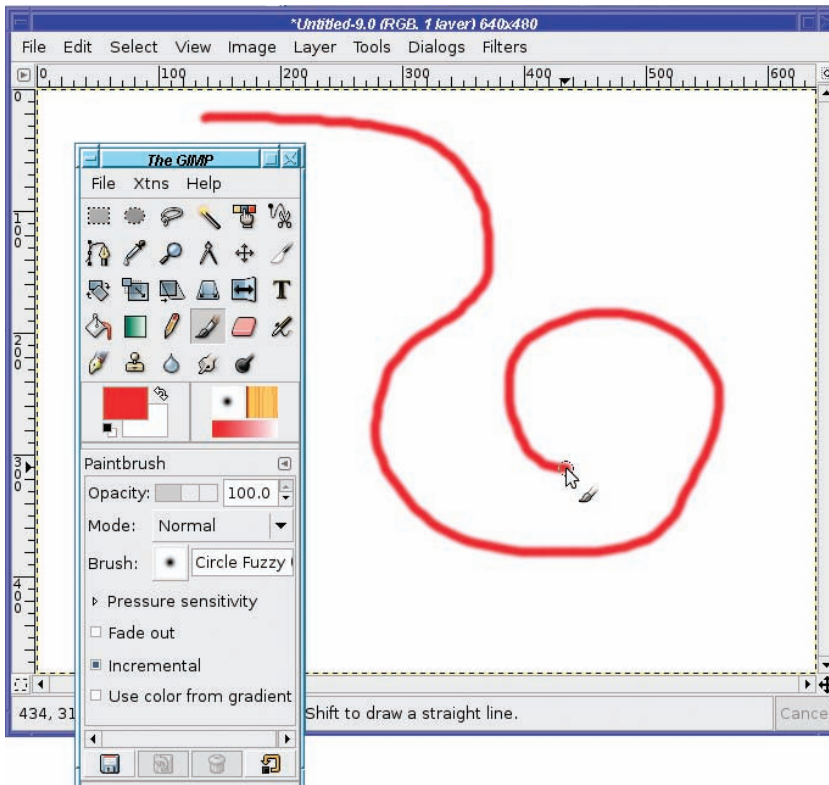
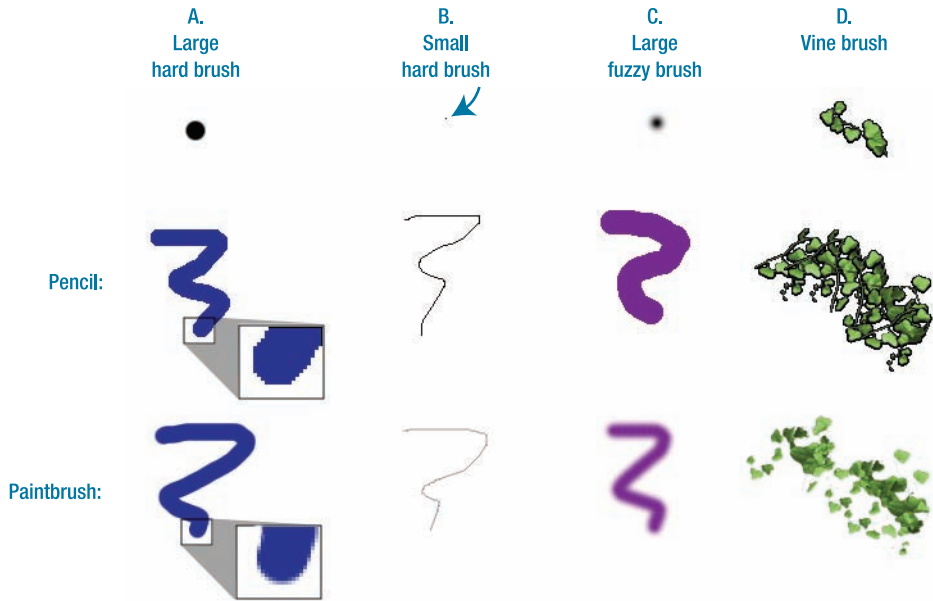


Figure 4-9. The Paintbrush tool

The Paintbrush tool differs from the Pencil tool in two important ways. First, it can use fuzzy-edged brushes. Second, it can use the hard-edged brushes too, but it uses them in a different way from the Pencil tool.

Figure 4-10 shows some of the differences. In 4-10A, each tool is used with a large, hard-edged brush. The results look similar until you look closely.



**Figure 4-10.** Differences between the Pencil and Paintbrush tools

The Paintbrush uses a technique known as *antialiasing* on the edges of diagonal lines: pixels along the edges are made semi-transparent, or blended into the background color, to fool the eye into seeing a smooth diagonal line. The Pencil tool does not use antialiasing, so the edges look jagged.

Then why would you ever want to use the Pencil tool? 4-10B shows you one reason. Antialiasing on a thin line can make it fade into the background.

The lower Paintbrush line in B was drawn in the same black color as the upper Pencil line. Notice the Pencil line is sharp, black, and distinct as compared to the fuzzy gray Paintbrush line. For small, fine artwork, the Pencil tool is often best.

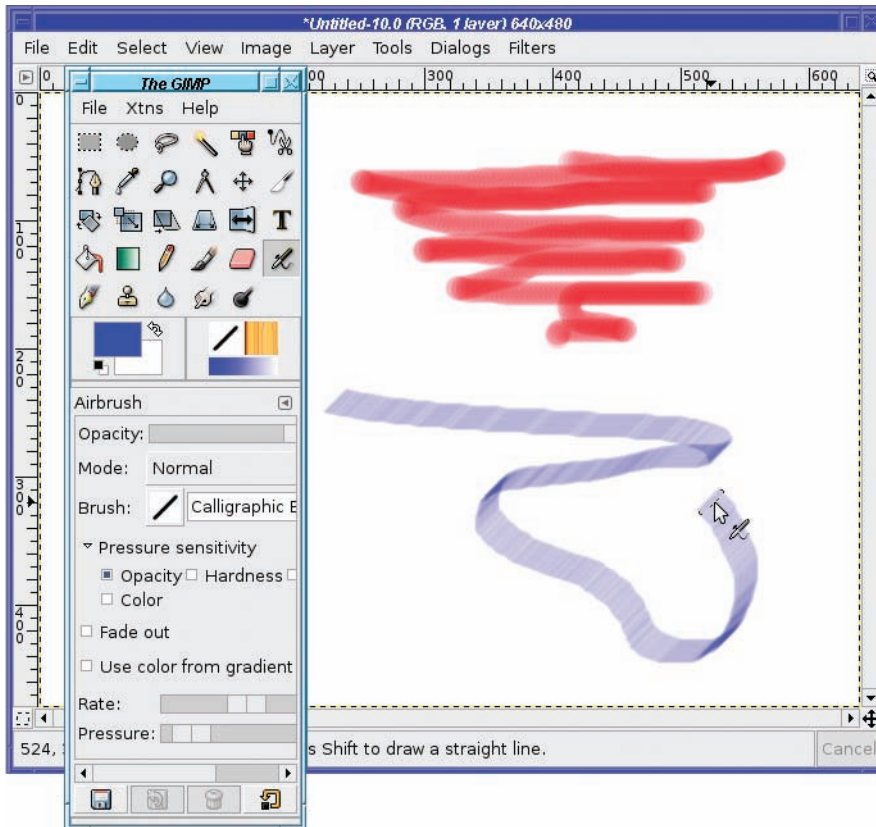
The other reason is indexed images. It takes more colors to draw an antialiased line. This means that the eventual file size will be larger. It also means the final image might not be usable for processes such as T-shirt or business-card printing, which can only handle a small fixed number of colors.

4-10C shows the tools when used with a fuzzy-edged brush. Obviously, the Paintbrush tool wins here. The Pencil tool ignores any fuzzy edges in the brush, and paints a wide, fat line.

4-10D shows how the choice of tool can make a difference with some of the more elaborate brushes. The Pencil preserves all the details of the vine leaves, while the Paintbrush creates an interesting sponge-art effect.

## The Airbrush Tool

The GIMP offers two more tools for drawing lines. One of them is the Airbrush tool (Figure 4-11).



**Figure 4-11.** The Airbrush tool, using a large, hard-edged brush (top) and the slanted calligraphic brush (below)

The Airbrush almost always draws fuzzy edges, even if you use a hard-edged brush. It's also time sensitive. The slower you drag across the page, the darker the line will be, just as if you were using a real airbrush or a can of spray paint. Hover over one spot holding the mouse button down, and the spot will get darker and darker until it's completely opaque. (That means that if you stay in one place with a hard-edged brush, eventually the edges will become sharp. That's how the Airbrush can draw non-fuzzy edges.)

Painting with the Airbrush takes some practice and a deft touch . . . just like a real airbrush! The Airbrush adds two additional options to the normal set:

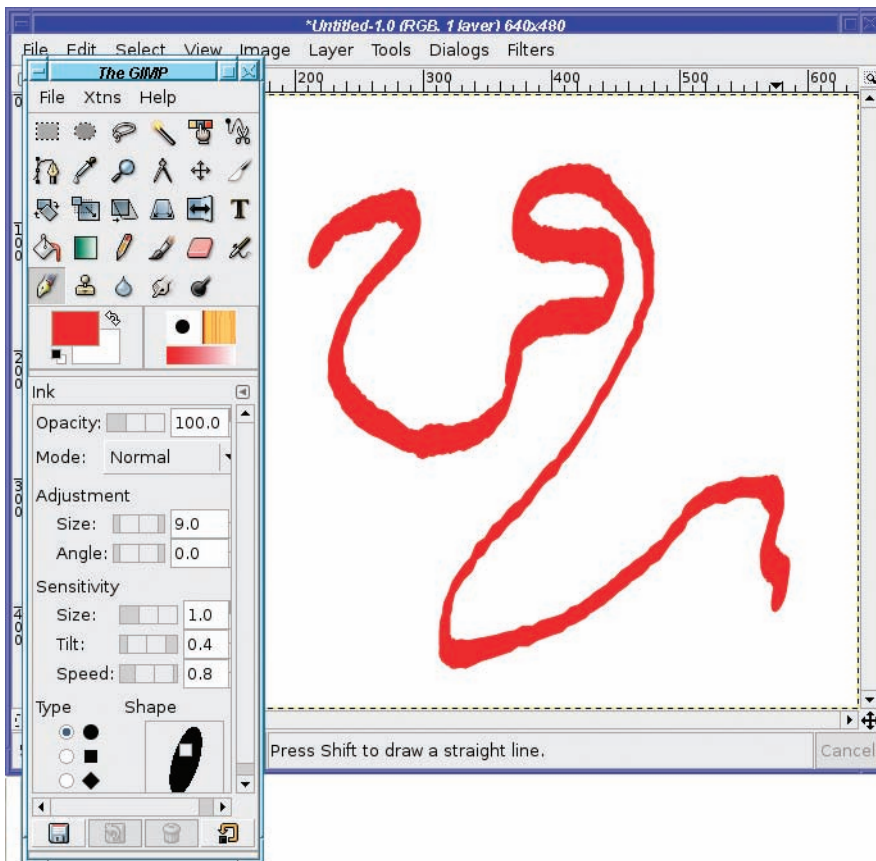
- *Rate* controls how sensitive the Airbrush tool is to movement, i.e., how fast things get darker if you slow down.
- *Pressure* controls the darkness. Think of it as the amount of paint the airbrush is spraying.

## The Ink Pen Tool

The Ink Pen tool (Figure 4-12) is the most complex of the basic drawing tools, but it can be very rewarding. It emulates an old-fashioned fountain pen with replaceable nibs (tips).

Not only can you make nibs of different shapes (analogous to choosing different brushes with the other tools), the Ink Pen reacts to your speed as you draw, creating a line that varies in width like a line from a real ink pen.

Careful: you can create blobs of ink if your penmanship isn't quite right. But unlike a real ink pen, you won't get ink all over your shirt pockets!




**Figure 4-12.** *The Ink Pen*

The Ink Pen is most useful if you have a drawing tablet rather than a mouse. It will change line thickness in response to pressure or tilt as reported by a tablet. But even with a regular mouse, you can sometimes create nice effects.

The Ink Pen, unlike all the other drawing tools, ignores GIMP's current brush. It has its own notion of a brush, similar to GIMP's parametric brushes, with the following options:

- *Adjustment* controls the width of the pen's nib and its angle (0 is horizontal).
- *Sensitivity* controls how much the Ink Pen will react as you vary your drawing style with your speed or the pressure and tilt on a tablet.
- *Type* controls the shape of the nib. Initially, you choose one of three basic shapes.
- The *Shape* box to the right of the type selector lets you refine the nib. It shows the chosen shape with a small square in the middle of it. Drag the square to change the aspect ratio and tilt of the nib (Figure 4-13).



**Figure 4-13.** Dragging the box to change the nib's shape

## USING A DRAWING TABLET

Drawing tablets, or graphics tablets, are devices that include a pad, a stylus, and sometimes a mouse or additional buttons. They're available in a variety of sizes from several manufacturers.

Tablets are very popular with professional computer artists for two reasons. First, drawing with a pen-shaped object feels more natural to most people. You have much more control over shapes, and most people find that they can make lines flow more smoothly with a pen on a tablet than with a mouse.

The other important reason? Most graphics tablets offer additional controls. They're sensitive to *pressure*: The GIMP can vary line width, color, or opacity of the line you draw depending on how hard you press on the tablet. Some tablets are also sensitive to *tilt*: The GIMP can detect whether you're holding the pen straight up or tilted over on its side, and can vary the line accordingly.

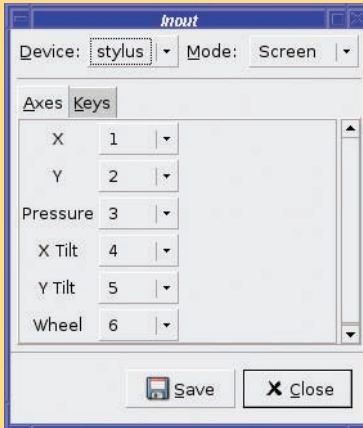
All of GIMP's drawing tools allow you to specify pressure sensitivity if you have a tablet, and some (like the Ink Pen) are sensitive to tilt as well.

Many tablets use a double-ended stylus: the upper end works as a second tool, usually called the *eraser*. The GIMP can assign different tools to the stylus and the eraser, which makes it very convenient to switch between the two simply by flipping the stylus over.

There's a catch to all this, though. It's sometimes difficult to get the operating system to see all the drawing-tablet information the GIMP needs.

Tablets under Windows usually work, though there may be some problems with particular models. Tablets under Linux often require you to download additional drivers for the kernel and for X. On the Mac, Apple's implementation of the X Window System does not currently support drawing tablets, though the free Darwin X server does.

If you have a drawing tablet, you can tell the GIMP about it by clicking on *Configure Extended Input Devices* in the *Input Devices* section of the Preferences window. This will pop up the Input dialog, as shown in the following image:

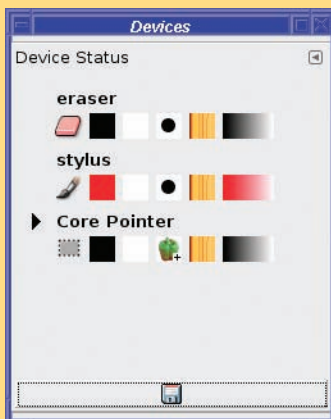


By default, the tablet's tools will have *Mode: Disabled*. Set this to *Screen* for each tool in turn, then click *Save*, and GIMP should remember the settings next time.

The other *Mode* option besides *Screen* is *Window*, which restricts the tool to work only in the image window. This in theory could give you better control, especially with the smallest-sized tablets. However, in practice, it doesn't work very well on most systems, and it makes it impossible to move the stylus back to the Toolbox window to choose a new tool.

Try *Window* if you want, but most people find that setting all tools to *Screen* works best.

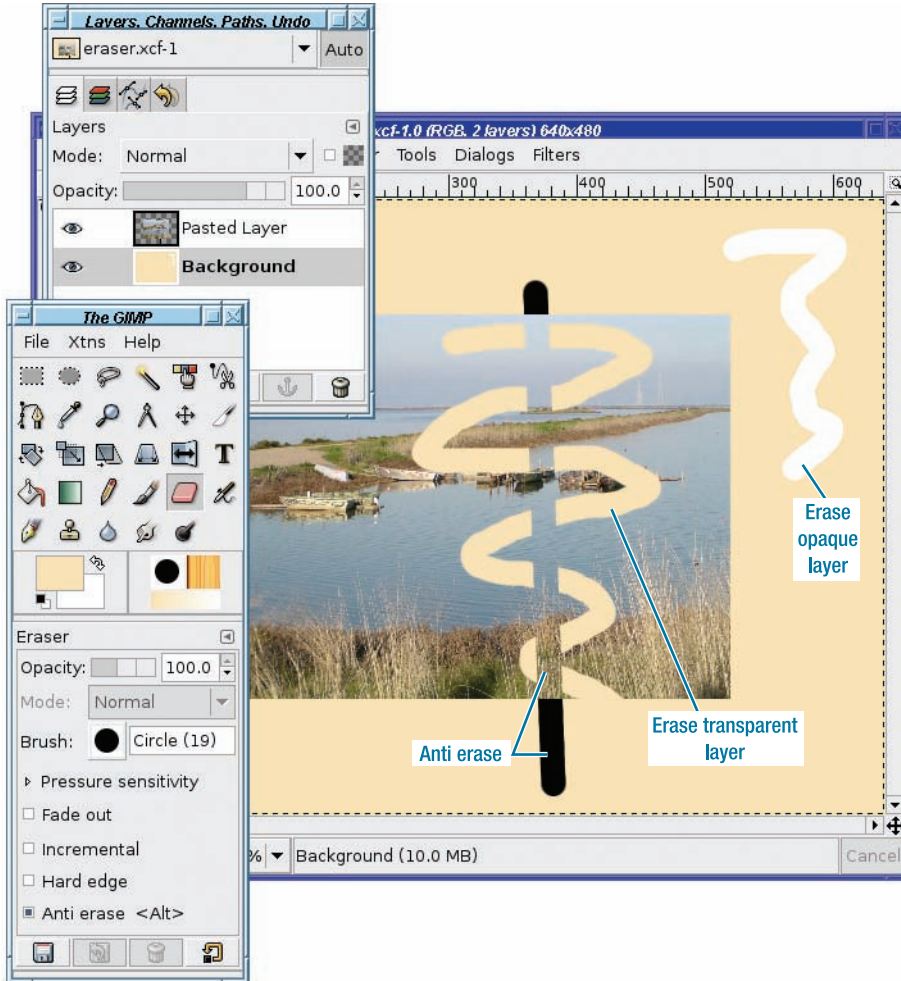
The GIMP also has a *Device Status* dialog, accessed from the Toolbox via *File > Dialogs > Device Status*, as shown in the following image. This shows the pointer devices GIMP currently knows about (including tablet devices and the mouse, denoted *Core Pointer*). It also shows each device's assigned tool, foreground and background color, brush, pattern, and gradient. You can drag and drop items in this dialog to change the values there.





## The Eraser Tool

The Eraser (Figure 4-14) erases. Simple, no?



**Figure 4-14.** *The Eraser tool*

Well, not *quite* that simple. The definition of *erase* depends on the image. If the current layer has an alpha channel, then erasing makes the current layer transparent wherever the eraser touches it.

Remember, *alpha channel* means that a layer or image supports transparency. The GIMP uses alpha as a synonym for transparency in quite a few places. Tutorials on the web for GIMP or other image-editing programs use this term quite a bit, so it's good to remember it.

In a layer that does not support transparency, the Eraser just paints the current background color.



The Eraser, like other drawing tools, uses the current brush. Use a large brush for erasing large areas. For erasing smaller areas or single pixels, use a smaller brush. Remember, you can zoom in on an image to magnify it so that you have more control over exactly what gets erased.

Most of the time, there are more precise ways to remove detail than the Eraser. (You'll learn about some of these methods in Chapters 5 and 6.) But the Eraser is useful for quick-and-dirty touch-ups of drawings.

The Eraser adds two new options to the usual drawing tool set. *Hard edge* makes the Eraser ignore any fuzzy edges in the brush. Remember how the Pencil tool made fuzzy-edged tools look large and sharp-edged? The *Hard edge* option will make the Eraser behave like that.

Anti erase is trickier. How is it possible that GIMP could un-erase something?

When you erase in a layer with an alpha channel, you're merely telling GIMP to make that part of the layer transparent. Behind the scenes, though, the GIMP knows which colors were there before you made it transparent. With the *Anti erase* option, GIMP removes transparency from the layer wherever you drag the eraser. If there was color there that you erased earlier, you'll see color back again.

---

**Caution** Using *Anti erase* on a transparent area that has always been transparent will leave a black trail, as though you were drawing with the Pencil with a black foreground color.

---

Figure 4-14 shows the difference in how the anti-eraser behaves in various circumstances. The image has two layers: an opaque background layer, and another layer that includes a photograph in the center, with transparency everywhere outside of the area of the photograph.

Selecting the opaque background layer and erasing leaves a trail of white (the current background color). Anti-erase on this layer does nothing, since there's no transparency to remove.

Selecting the pasted photograph layer gives more complex behavior. Erasing over the photograph makes the photograph transparent, so you can see through to the background layer. Notice that where the erased trail moves off the bottom of the photograph into the area that's already transparent, nothing changes: erasing transparency just gives you transparency.

Anti-erase over the photograph does nothing, except in areas that were previously erased. There, it un-erases, removing the transparency and revealing the photograph again. But where anti-erase strays outside of the photograph into purely transparent areas, it erases the transparency and replaces it with black, since there was previously no other color there.

## Drawing Rectangles, Circles, and Other Shapes

Now that you're an expert on drawing lines and erasing them, you'd probably like to draw some shapes, too.

### Defining Regions

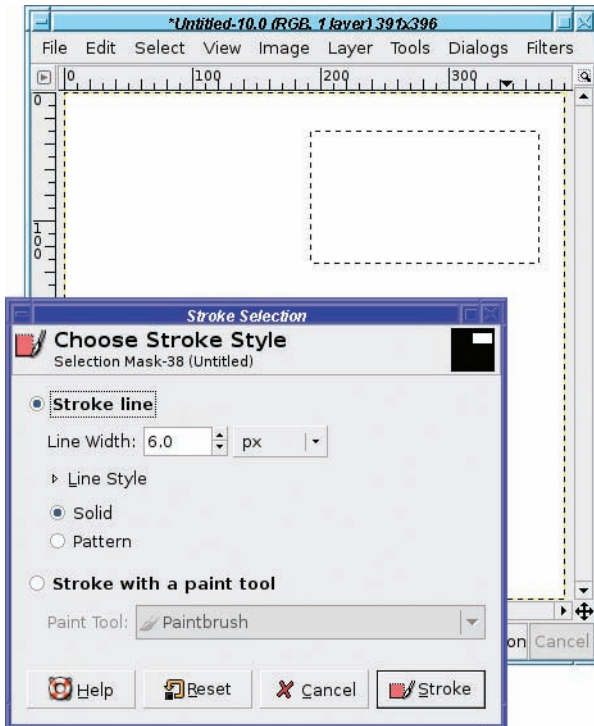
The GIMP doesn't have any tools specifically for drawing shapes, which confuses some people coming to the GIMP from vector-drawing programs. But GIMP doesn't need shape-drawing tools: the trick is to make a selection, then *fill* or *stroke* the selection.

To try this, first use the Rectangle or Ellipse Select tools and drag out a selection in the image.

**Tip** To constrain the Rectangle Select tool to a square, or the Ellipse Select tool to a circle, press the Shift key *after* you start dragging in the image. Pressing Shift first might sometimes work, but often it will give surprising results, which will be explained in Chapter 5.

## Outlining Selections: Stroking

*Stroking* a selection outlines the boundaries of the selection. It's located in *Edit* ► *Stroke Selection...*, and brings up the Stroke Selection dialog (Figure 4-15).



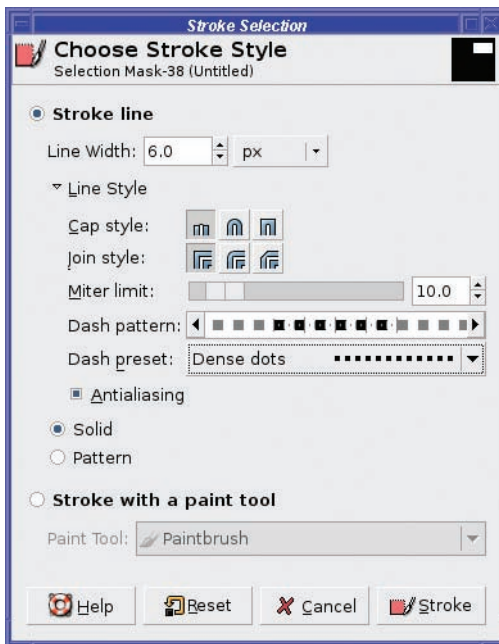
**Figure 4-15.** The Stroke Selection dialog, with a rectangular selection active

*Stroke line* will draw outlines using the current foreground color. You can specify the *Line Width* in pixels (or, if you choose, in some other unit). The default 6-pixel line is quite wide.

Alternatively, you can stroke with a line-painting tool, such as Pencil or Paintbrush. The GIMP will act as though you had run that paint tool around the borders of the selection, using the current color and brush. (You can even make picture frames by stroking with a wide brush, especially if you use one of the more elaborate brushes, such as some of the animated ones.)

Sometimes stroking with a paint tool, especially an antialiased tool like the Paintbrush, works better on curves than specifying a line width. Try it both ways on a circular selection and see.

If you use *Line Width* rather than a paint tool, you can specify the type of line that will be used by clicking the expander next to *Line Style* (Figure 4-16).



**Figure 4-16.** *Line styles in Stroke Selection*

*Cap style* changes the look of the ends of lines; *Join style* does the same for places where lines connect to other lines (for instance, at the corners of a rectangle).

*Miter limit* specifies how *sharp* or *blunt* the corners will be. A limit of zero makes the corners as sharp as possible; higher numbers make them look more blunted or rounded. Cap and join styles and miter limits won't be very noticeable unless your line is quite wide.

*Dash pattern* lets you make a dashed line and specify exactly how the dashes will appear (drag horizontally in the preview to make different patterns). *Dash preset* lets you choose from a selection of common dashed-line styles.

You've already seen what *Antialiasing* can do, in the section on the Paintbrush tool. It can be helpful on rounded or diagonal lines.

You can also stroke with a pattern (which generally will only be interesting if you use a fairly large line width). I'll talk about patterns in a moment.

---

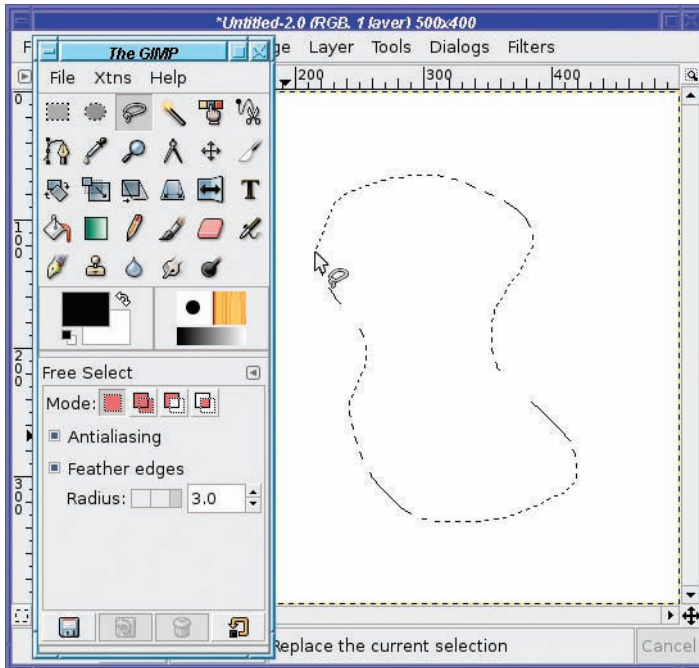
**Note** Both rectangular and elliptical selections offer *antialiasing*, which uses transparency along the edges of the selection to make boundaries appear smoother. They also let you *feather edges*: you can make the edges of your selection fuzzy rather than sharp. This isn't useful for outlining selections, but it can be useful when you're filling them. There are also some other selection options that you'll learn about in Chapter 5.

---

## Free Select with the Lasso

Another selection tool you'll find useful for drawing is Free Select, also called the Lasso tool, from its Toolbox icon.

The Lasso lets you draw a freehand selection of any shape (Figure 4-17). Just outline your shape in the image, end somewhere near where you started, and GIMP will make a closed curve. You can treat it just like you would a rectangular or elliptical selection.

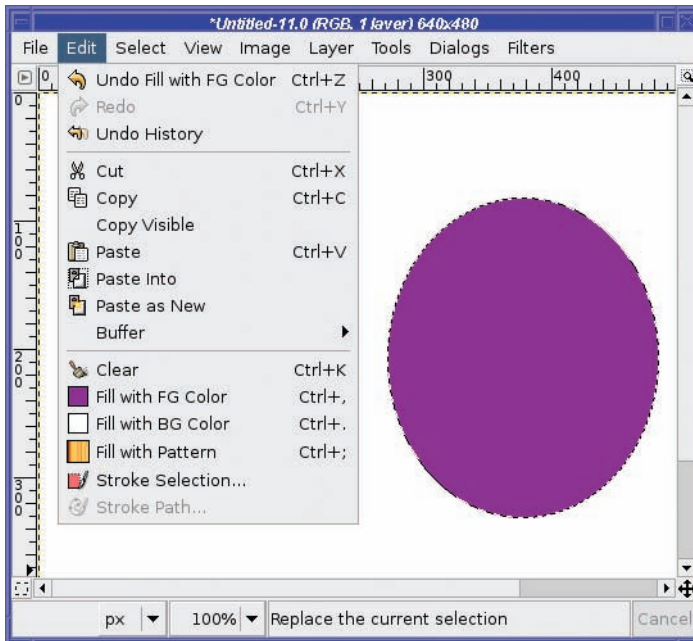


**Figure 4-17.** *The Free Select tool*

Free Select offers the same antialiasing and feather options as the other selection tools.

## Filling Regions

Filling a selection with a solid color or pattern is even easier than stroking a selection. Just above *Stroke Selection* in the *Edit* menu, *Edit* ► *Fill with FG Color* (Figure 4-18) does the obvious: it fills your selection with the current foreground color. *Fill with BG Color* does the same thing but uses the background color, and *Fill with Pattern* fills the selection with the current pattern.



**Figure 4-18.** *Fill with foreground color, applied to a selection*

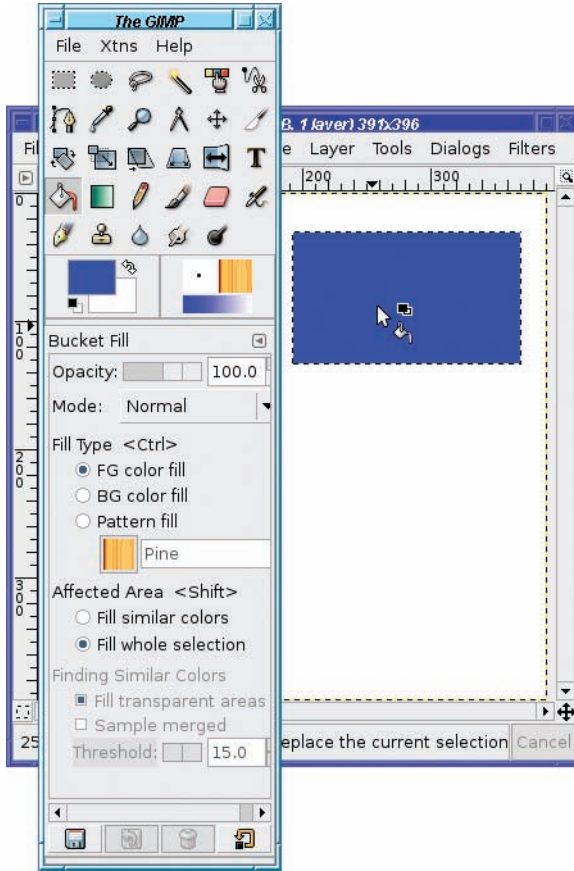
Notice that all three of these *Fill with* items show the current color or pattern next to the entry in the menu. This gives you a quick sanity check to make sure you're using the color you expect. Of course, if you decide that you don't like the color, just go to the Toolbox and click on the color swatch to change the color, then *Fill with FG Color* again.

If you don't have a selection, the *Fill with* items will fill the whole layer. This is an easy way to change the color of a background layer.

You can also fill a selection by dragging from the Toolbox's foreground, background, or pattern swatches into the image.

## The Bucket Fill Tool

For more complicated fill-ups, use the Bucket Fill tool (Figure 4-19). Choose it in the Toolbox, then click anywhere inside your selection to fill it with the foreground color.



**Figure 4-19.** *The Bucket Fill tool*


Bucket Fill has lots of options, starting with the *Opacity* and *Mode* selections shared with the line-drawing tools.

*Fill Type* lets you fill with the foreground color (FG), background color (BG), or a pattern. The <Ctrl> next to *Fill Type* tells you that pressing the Control key when you click in the image will choose the opposite of your specified color: if FG is chosen, then Control-click will fill with the background color rather than the foreground, and vice versa.

*Affected Area* lets you control what area will be filled. *Fill whole selection* is easy: select the area you want to fill, and click in it. But with *Fill similar colors*, when you click somewhere in the image, GIMP will look for adjacent areas that are colored similarly. (If there's a selection, GIMP still won't fill outside it.)

Just *how* similar the color needs to be is controlled by the *Threshold* slider (Figure 4-20).





**Figure 4-20.** *Fill similar colors*

For example, in a photo like Figure 4-20, you might want to use *Fill similar colors* to replace a cloudy sky with a blue one. Set the foreground color to light blue, then click somewhere in the sky.

With a very low *Threshold*, not much of the sky is likely to have the same color as the point where you click; you end up with color bands and an unnatural look. With the default *Threshold* of 15, a lot of the clouds get replaced, but not all of them; sometimes this looks okay, but other times it looks weird. With just the right *Threshold* (50, for this image), the whole sky is replaced, but nothing else. Go higher, and some of the fill will spill over into other areas, like the cliff face. At very high settings, nearly every color in the image will be replaced, and at the highest setting of 255, the entire image will be blue. The only way to determine the exact *Threshold* setting is by experimenting; it's different for every image.



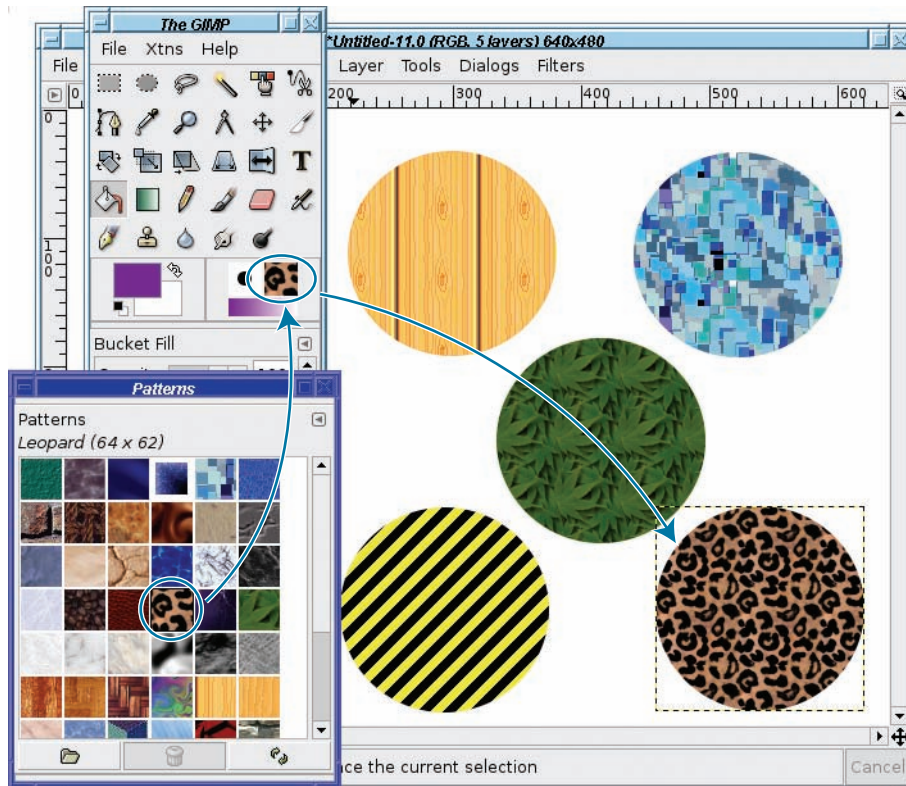
*Fill transparent areas* lets you fill similar colors, even if you click on a transparent or nearly transparent area. Normally, clicking on transparency won't fill anything.

*Sample merged* uses the colors from all layers in the image to determine the color GIMP will try to match, instead of just the color in the current layer. This is useful when you're building up an image with a lot of layers stacked on top of one another. Of course, when GIMP fills, it will still put the color in the current layer.

## Patterns

What's all this about patterns?

The GIMP has a diverse collection of patterns (Figure 4-21). The current pattern is always shown in the Toolbox, next to the current brush. It will also be shown next to tools and menu items that use patterns, such as Bucket Fill and *Edit* ► *Fill with Pattern*.



**Figure 4-21.** Filling circular selections with patterns

Clicking on the pattern shown in the Toolbox will bring up the Patterns dialog, where you can select from all the patterns GIMP knows about.

A pattern is just an image. Right-clicking on any square in the Patterns dialog offers *Open Pattern as Image*, from which you can edit the image and save it as a new pattern (usually with a file name ending in .pat, though .png and .tga will also work).

But a pattern is a slightly special type of image: it's usually *tileable*. Imagine the pattern image is a tile and you're flooring a room with it. You'd want the right edge of each tile to match the left edge of the tile next to it, and similarly with the top and bottom of each tile.

If the edges match, then the tile job will be seamless and look like a single huge pattern. If they don't match, then you will be able to see seams between each image. It will be obvious that you used a lot of small images, and the result won't look so good.

Most of the time, the patterns built into GIMP offer whatever you need. If you want more, there are plenty of sites that offer downloadable patterns. But if you crave something really special, the GIMP can help you make an image into a tileable pattern. *Filters* ► *Map* ► *Make Seamless* changes the edges of the current image so that it can be tiled precisely. Chapter 9 will go into pattern creation in more detail.

## Gradients

The Gradient, or Blend, tool is another way to fill areas (Figure 4-22). It's also useful in all sorts of GIMP effects, as you'll see in later chapters.




Figure 4-22. Filling with a gradient

The typical gradient makes a smooth transition from the foreground color to the background color. (There are also other gradients that display different effects, such as *Golden* for making metallic golden glows, or *Neon Cyan* for making neon lights.) With the Gradient tool selected in the Toolbox, mouse down in one place, drag to another place, and release the mouse button to draw a gradient into the current layer (of course, if there's a selection, the tool will only draw inside it).

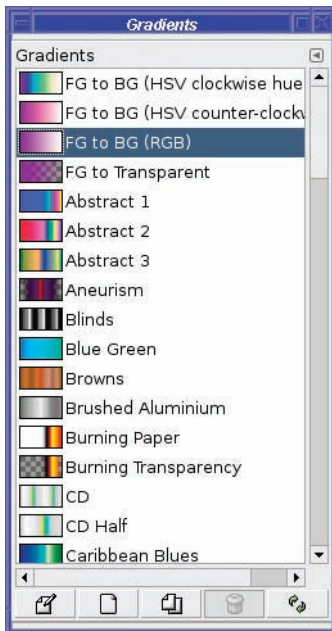
---

**Tip** If you drag even slightly off horizontally or vertically, you'll get a gradient that's noticeably diagonal. If that isn't what you want, try holding down the Control key as you drag. This will constrain your dragged line to a multiple of 15 degrees, which will make it much easier to drag a perfectly horizontal or vertical line. This Control-key constraint works in most drawing tools, and also in the transform tools such as Rotation.

---

## The Gradients Chooser

You may have already noticed that there's a gradient indicator in the Toolbox, below the current brush and pattern. That's because the default gradient, a smooth blend from foreground to background color, isn't the only gradient available (Figure 4-23).



**Figure 4-23.** *The gradients chooser*

Most of the gradients available in the dialog have fixed colors regardless of the current foreground and background. Some gradients also include transparency, or partial transparency. You may find that some of them give a nice effect, or you may find most of them extraneous.

Your call! But the basic gradient, *FG to BG (RGB)*, is useful for all kinds of effects, as you'll see in later chapters.

The Gradients dialog has buttons at the bottom similar to the ones in the Brushes dialog: *Edit*, *New*, *Duplicate*, *Delete*, and *Refresh*. As with Brushes, you can create new gradients and save them to use later.

## Other Gradient Options

The usual drawing tool options, *Opacity* and *Mode*, are available in the Gradient tool options you saw back in Figure 4-22.

*Mode* is more interesting here than in most of the other tools, as it can create some wild effects. For instance, try using *Difference* mode and applying the same gradient repeatedly (with a colored gradient; it's not very interesting in black and white). Then try it on a gradient of a different shape (see *Shape*, described in this section).

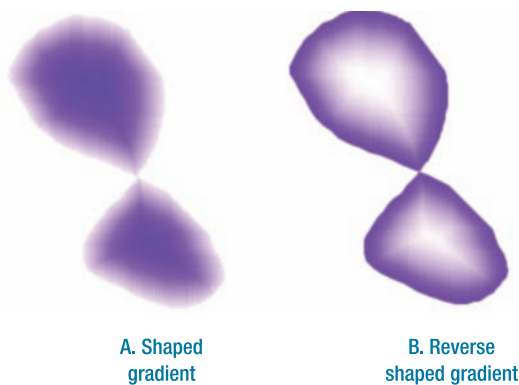
What *Mode* actually does will be explored in more detail in Chapter 9. For now, just play around and see what you can make.

*Reverse* reverses the colors of the gradient. Instead of starting with the foreground color and blending into the background color, it will start with the background color. Of course, you could just drag in the other direction to get the same effect; but some of the other gradient shapes only draw in one direction. With them you *will* need the *Reverse* button.

*Offset* controls where the gradient actually starts, as a percentage of the total distance you drag. If it's 0 (the default), then GIMP will create a gradient from your start point to your end point. If *Offset* is 40, then GIMP will draw the foreground color as a solid color for the first 40% of the distance, and only then will start the gradient. You can accomplish the same effect by simply starting your drag in a different place, so this option isn't usually needed.

*Shape* offers a wealth of other ways GIMP can orient the gradient. *Linear* is the default, but GIMP can also make gradients that look like circles, cones, spirals, or (the most fun) it can make a gradient following any shaped selection you make.

Try making a shaped selection using the Lasso tool; then set the gradient to *Shaped (Spherical)* and drag inside the image (Figure 4-24A). Then change the gradient to *Reverse* and drag again (Figure 4-24B).



**Figure 4-24.** Shaped gradients (spherical) on a free selection

Even if you drag over only a small distance, the Gradient tool will fill the whole selection (or image, if there's no selection active). *Repeat* gives you control over what happens outside of the region you dragged. The default setting, *None*, will simply use the beginning and ending colors of the gradient. If you select *Sawtooth wave*, GIMP will repeat the gradient, starting again at the beginning color; with *Triangle wave*, the gradient will be repeated in reverse.

*Dithering* offers higher precision in the way the gradient's color is calculated. Most of the time you won't notice a difference, but under some circumstances, you might see color-banding with this option disabled, especially on very large images.

*Adaptive supersampling* is another way of getting higher precision, at the cost of more computational work. Most of the time you probably don't need it, but if you notice unevenness or a stair-step effect, try it.

### WHEN THE GIMP WON'T DRAW

The GIMP is a great drawing program, but it can be frustrating too. Sometimes, you pick a drawing tool, try to draw, and nothing happens. What's up?

There are a lot of reasons that this can happen. Here are some of them:

- If you have a selection, GIMP will only draw inside it. When you try to draw outside the selection, nothing happens. Usually, it's obvious when you have a selection, because you can see the “marching ants” selection outline. But it's possible to have a selection and not know it. You might have done *View* ► *Toggle Selection* at some point to hide the selection boundaries, and then forgotten to turn them back on. The selection might be there but too small to see. Or, it might be scrolled off the visible part of the page, if your image is large or you're zoomed way in. When the GIMP won't draw, try a *Select* ► *None* to make *sure* you don't have a selection in the image.
- The GIMP will only draw to the current layer. Do you have the right layer selected in the Layers dialog? Is the layer visible? Is the opacity set to something more than 0%, so that the layer isn't completely transparent? Do the layer's boundaries extend to where you're trying to draw?
- It's also possible you've pasted recently and still have a floating selection. GIMP won't draw normally to a floating selection. You'll have to click the *New* button to make it a new layer, or the *Anchor* button to merge it into the previous active layer.
- While you're looking at the Layers dialog, check the *Keep Transparency* button on the layer (the one to the right of *Mode*, with the checkerboard square next to it). If that's checked, then GIMP won't let you draw on currently transparent parts of the layer.
- Check the opacity of the tool in its tool options: make sure it's at 100%.
- Layers each have a mode, and some of the modes do strange things in combination with the layer beneath. Make sure the current layer's mode is *Normal* (unless you're deliberately creating an effect using a different layer mode, of course).
- Check the *Mode* in the tool options for the drawing tool you're using. If it's not *Normal*, unexpected things can happen. And usually will.

## A Drawing Project

It's time to put all these techniques together and create a drawing.

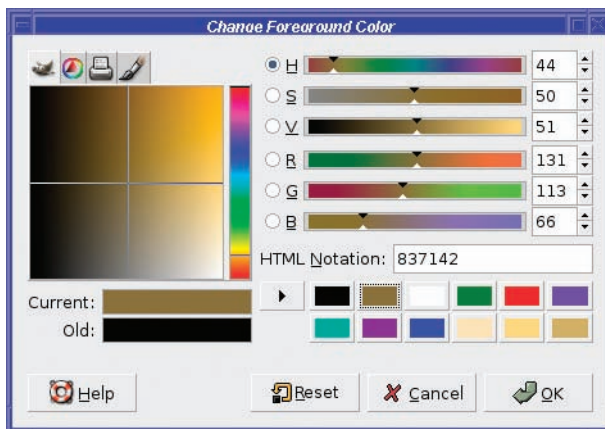
So bring up a blank canvas, whatever size you like, with *File* ► *New*. If you already have a canvas you've been scribbling in, *Edit* ► *Clear* (the Delete key in GIMP 2.3, Control+K in earlier versions) will clear out everything and leave you with an empty canvas.

### Drawing a Tree

For a first drawing project, how about a tree? A tree is fairly simple and doesn't require a lot of art experience.

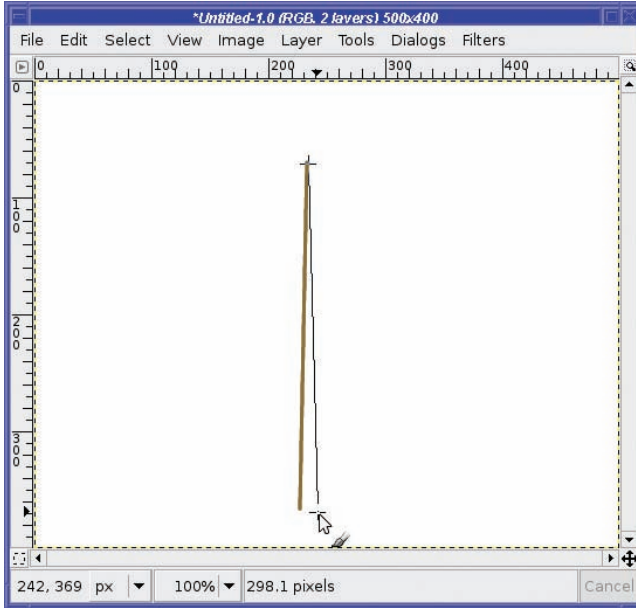
What's the first rule of drawing?

Right—create a new layer. Name it “trunk.” Choose a brown foreground color. It's not always obvious where to find brown in a computer's color chooser. Try setting the hue to orange, then drag left and down to make the color darker (Figure 4-25).



**Figure 4-25.** *Where to find brown*

You can make the trunk a straight vertical line, if you like. But a real tree has a trunk that tapers a little, narrow at the top and wide at the base. You can accomplish this by using a narrow brush and drawing several near-vertical lines. Click where you want the top of the tree to be, and then click to the left side of the base and Shift-click. Now click again at the top, and Shift-click at the bottom-right (Figure 4-26).



**Figure 4-26.** *The outline of the trunk*

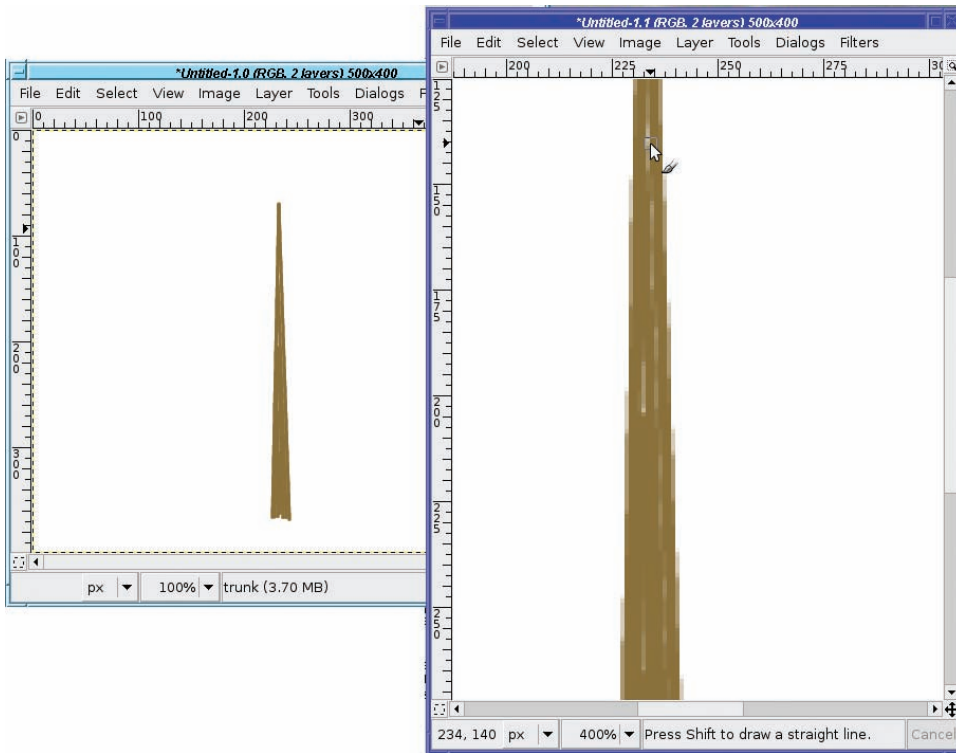
Then fill in the trunk by drawing more lines just inside your original outlines until the trunk is a fairly solid brown color. Or just until you like it!

If you have trouble filling in small holes in the trunk, you can zoom in, but there's an even better approach: multiple views. *View* ► *New View* creates another window showing the same image. Any changes you make in one window will be reflected in the other. The advantage? One window can be zoomed way in to see fine detail while the other window is zoomed out so you can see the results at actual size (Figure 4-27).

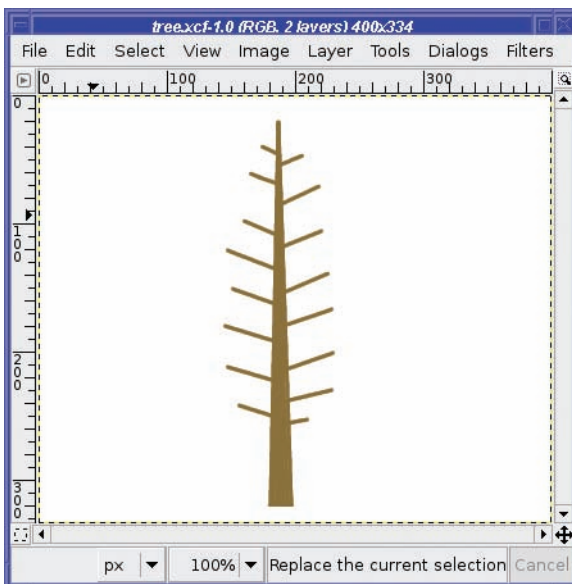
Once your trunk is ready, add some branches. Make a new "branch" layer. Then use the same click, Shift-click technique with the Paintbrush (Figure 4-28).

Now it's time to add the leaves (or needles, if it's a pine tree), so switch to a green foreground color. Use a new layer for this! You may decide later that you don't like your leaves and want to change them. Or you might want to make an autumn or winter version of your tree.



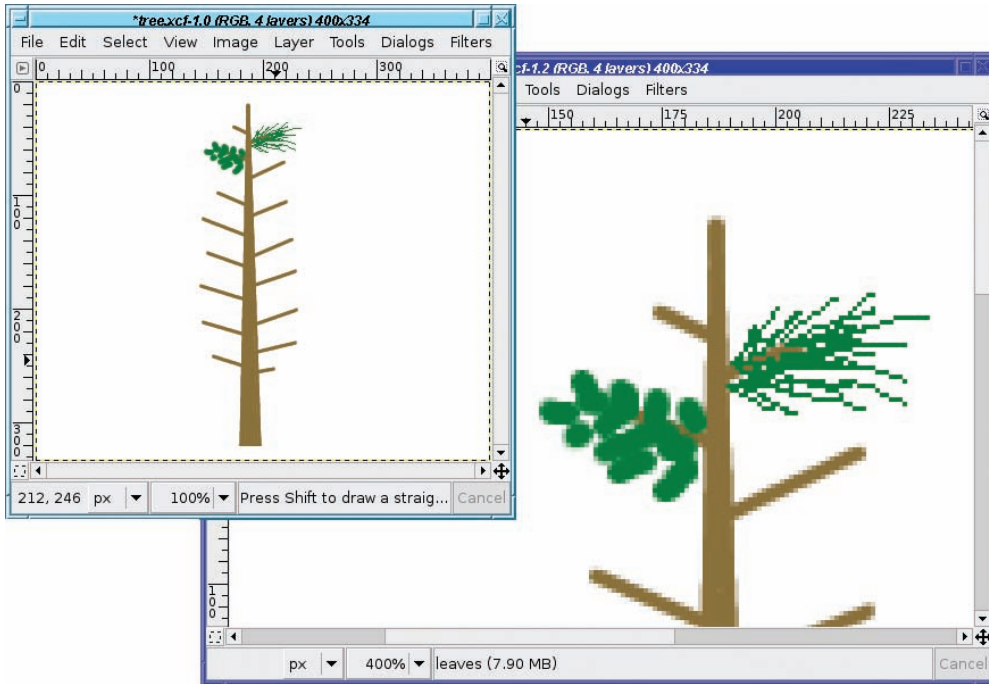


**Figure 4-27.** Using two simultaneous views



**Figure 4-28.** Add branches in a separate layer

For pine needles, use the Pencil tool and the smallest brush available, since you'll want to make thin lines. (Remember, the Paintbrush doesn't work well on the smallest brush. Making your drawing quite a bit larger would work around this and would give smoother lines than the Pencil tool.) For leaves, you might want a slightly larger brush (Figure 4-29). Experiment and see what looks right for your tree.



**Figure 4-29.** *Different types of leaves*

Continue putting in leaves, using short strokes, until your tree looks fairly well filled-in (Figure 4-30). Add more branches (switching to the “branches” layer, of course) if you need to.

Real trees have a range of leaf colors. You may want to add some slightly darker leaves, and some slightly lighter ones. Use separate layers for each color, in case the blend of colors doesn't work out and you need to change it (Figure 4-31).

The tree looks pretty good now. But that solid brown trunk doesn't look right. It needs a shadow!

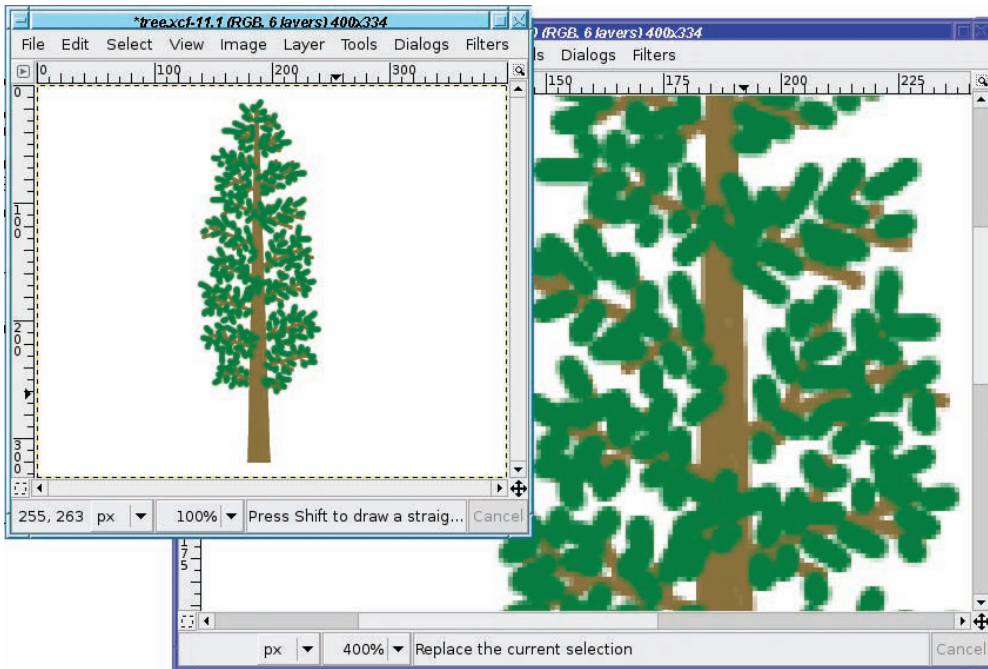


Figure 4-30. Leaves filled in

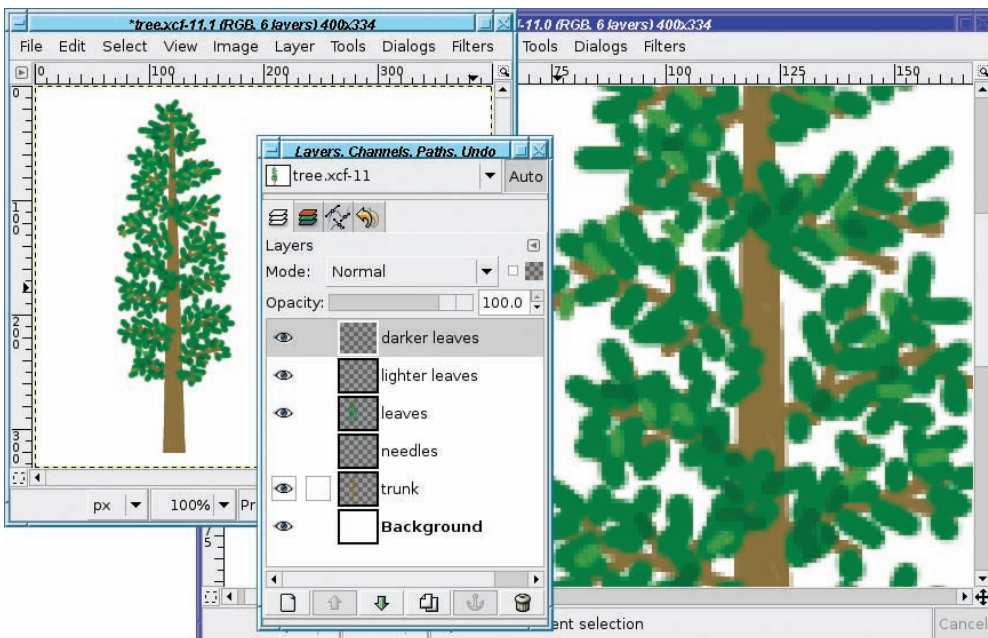


Figure 4-31. Try using several different leaf colors

To make a shadow on the trunk, turn off visibility on all the branch and leaf layers so you can see what you're doing. Select the “trunk” layer and duplicate it to create a new shadow layer called “trunk shadow,” just above the trunk layer but below the branches. Make this layer semi-transparent: try starting at about 35% opacity. You can adjust this later.

Select the Paintbrush tool (if it isn't already selected), a medium-sized fuzzy brush, and a dark foreground color. You can use black, but using a different dark color can give your shadow “depth,” as though it's reflecting light from objects nearby.

You're going to draw a dark shadow line up the right side of the tree. But first, a tip: you can make sure the shadow doesn't spill off the side of the tree by using the *Keep transparency* toggle button in the Layers dialog.

With *Keep transparency* checked, draw a line up the side of the tree trunk (Figure 4-32). Keep the center of the fuzzy brush about at the edge of the trunk—you want it to be blackest at the right edge of the tree trunk, and fade away toward the middle. It's okay that the big fuzzy brush hangs off the edge of the trunk; the *Keep transparency* option will prevent any overhang.

You can add shadows to your leaves, too. Since they're too small and numerous for you to draw all the individual leaf shadows like you did with the trunk, you can take a shortcut: make a drop shadow of the leaf layer or layers, using small values for offset and blur radius.

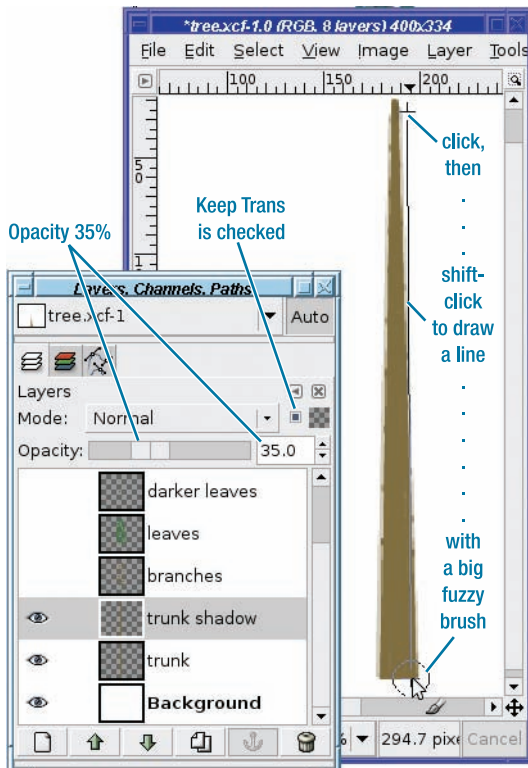


Figure 4-32. Drawing the shadow

Figure 4-33 shows the effect of the shadow. The tree is no longer flat—now it's a three-dee tree!



**Figure 4-33.** *The completed tree*

## Making a Box Using Perspective Transformation

It's nice to have a tree. But wouldn't it be helpful to have a wooden planter box so you could carry the tree around or take it inside?

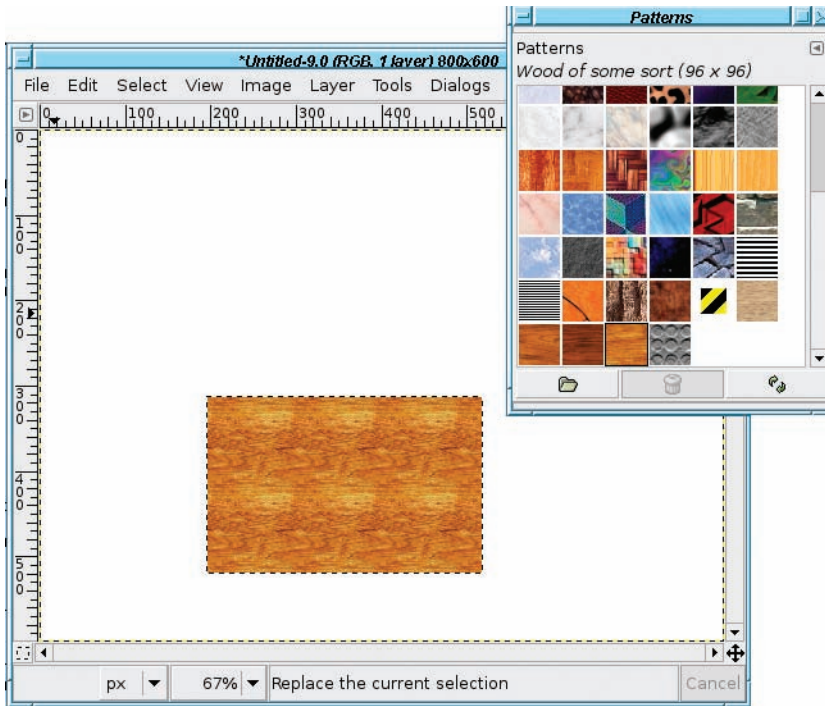
Make a new image to build your box. You can put the tree and the box together later. As with the tree, it's best to start with a blank image with a white background, and then make new transparent layers for each new piece you add. You'll want separate layers for each of the four sides of the box, so start with a layer called "box front" for the front face.

A box has a rectangular front. So make a selection using the Rectangle Select tool. You can fill that rectangle with a woodgrain pattern.

GIMP offers several woodgrain patterns to choose from. Some of them are better suited to a wooden box than others: the ones that look like wood paneling or a parquet floor look somewhat strange if you try to make a box out of them.

Just try lots of patterns until you find one you like. You can keep the Patterns dialog up while you choose a pattern, fill inside the image with *Edit* ► *Fill with Pattern* or Bucket Fill, and then choose another pattern. (If you undo between each step, you won't fill up your Undo stack with all the patterns you tried.)

I picked the birch-like pattern in Figure 4-34, which GIMP rather vaguely names “Wood of some sort.”



**Figure 4-34.** The first step in making a wooden box: a rectangular selection filled with a woodgrain pattern

Now, the box needs sides. You could try to draw a side using the Free Select tool. But really, the side of the box should be just like the front, except rotated so that it looks like it’s receding from you. GIMP has just the tool for that: the Perspective Transform tool.

First you need another layer. Make a duplicate layer of the box front, from the button in the Layers dialog. GIMP will name it “box front copy”; you’ll want to rename this to something like “box right side.”

The GIMP’s perspective transformations work best if the layer is no bigger than the object you’re trying to transform. (The image can be any size; it’s only the size of the layer that matters.) You can crop the layer to eliminate all the extra transparent areas using *Layer* ► *Autocrop Layer* from the image window menus.

---

**Note** If you still have your rectangular selection now filled with your pattern, another way to make a new layer containing only the selection is to copy, paste, and click *New Layer* in the Layers dialog to turn the floating selection into a normal one.

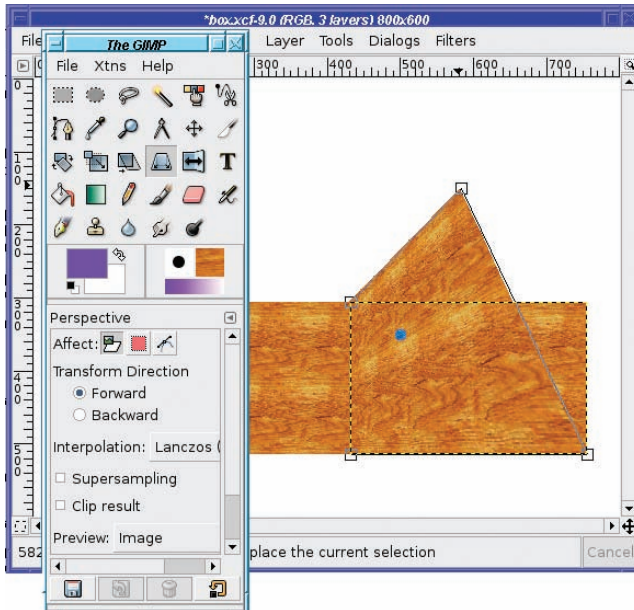
---



Use the Move tool to move the “box right side” layer over to the right, so that its left side touches the right side of the “box front” layer. It’s not critical that they line up exactly; just get them close.

Now select the Perspective Transform tool and click somewhere near the upper right of the box’s right layer (Figure 4-35). Drag to where you think the far top-right corner of the box will be.

Meanwhile, the *Perspective Transform Information* dialog pops up, giving you a bunch of data you probably don’t need, as well as buttons for *Help*, *Reset*, *Cancel*, and *Transform*. The last button is the one that will actually save the operation once you’re happy with it.



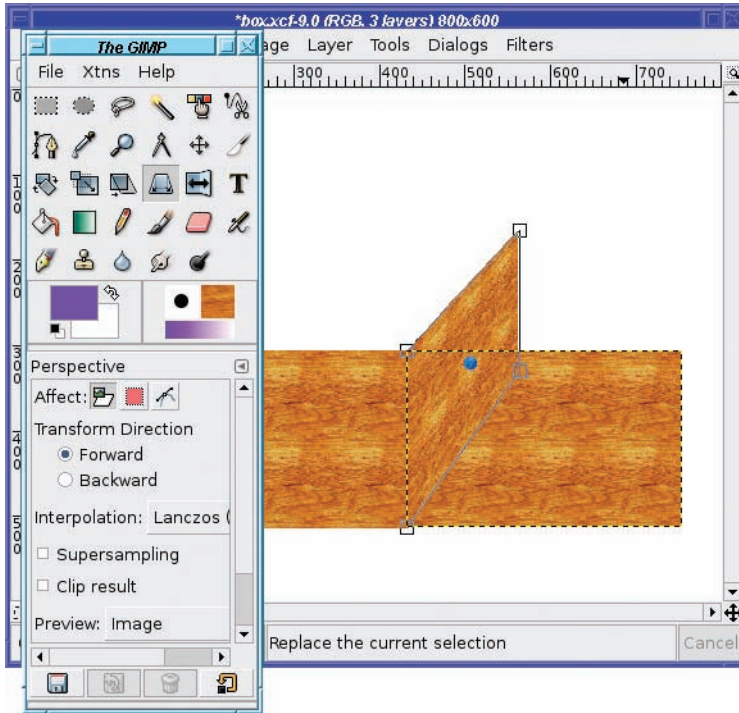
**Figure 4-35.** Beginning a perspective transformation

The Perspective Transform tool has roughly the same options as the Rotate tool you already used in Chapter 2. Most important, you can preview the actual *Image*, *Grid*, *Image+Grid*, or *Outline*.

Notice that the Perspective Transform tool gives you boxes at each of the four corners of the layer. These are *grab handles*. When you clicked near the upper-right corner of the layer, GIMP assumed you wanted that grab handle, so that’s the one that was dragged along with your mouse. (The blue dot shows you where the center of the layer will end up. It won’t necessarily be blue; the GIMP will show it in a color that contrasts with the color already there.)

Now click near the lower-right corner’s grab handle, and drag that up until it looks like it’s in about the right place for the box side (Figure 4-36).





**Figure 4-36.** *Dragging the second grab handle*

For a box like this, you want the far grab handles directly on top of each other: the line between them should be straight up and down, and should contain no “jaggies.” However, the diagonal lines—the “top” and “bottom” of the box as they recede—will not look right if they’re parallel: they should be closer together at the far end, just as the sides of a railroad track look closer together where the track is further away. (See Chapter 9 for a detailed discussion of perspective drawing.)

When you’re happy with the perspective, click *Transform* in the dialog.

You may find that it doesn’t look quite as good after you’ve clicked *Transform* as when you were previewing. If so, just undo and try again.

The next step is to make the left side of the box, pretty much the same way as you made the right side. Start by copying the box front. Crop the layer if necessary, then use perspective transform.

Bonus! On the left side you don’t need to move the layer, because the *left* side of the box left layer should touch the *left* side of the box front. That will happen automatically as soon as you copy the layer.

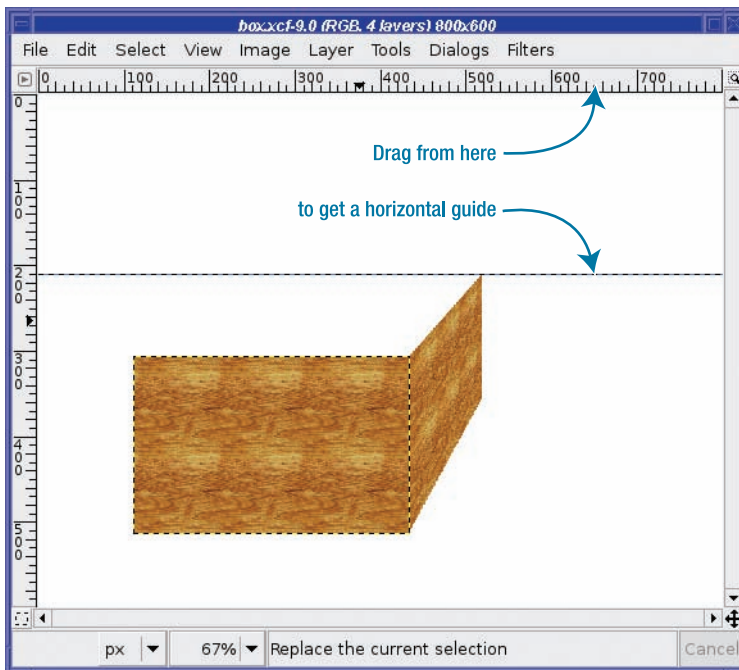
You might think you could just copy the box's right side layer to make the box's left side layer, but that won't look quite right. The sides of the boxes need to appear closer together at the far end than they are at the near end (they need to converge) in order for the perspective to look right. If you make the two sides exact copies of each other, it'll look wrong. (Don't take my word on this. Try it and see!)

## Using Guides

To make the box look right, it's best if the far top ends of the two box sides line up at the same height. A *guide* can help you with that.

Guides are horizontal or vertical lines you can put anywhere in the image to help you draw in exactly the right place. Guides don't show up in the final image; they're strictly there to help you draw, and you can *Show* or *Hide Guides* from the *View* menu.

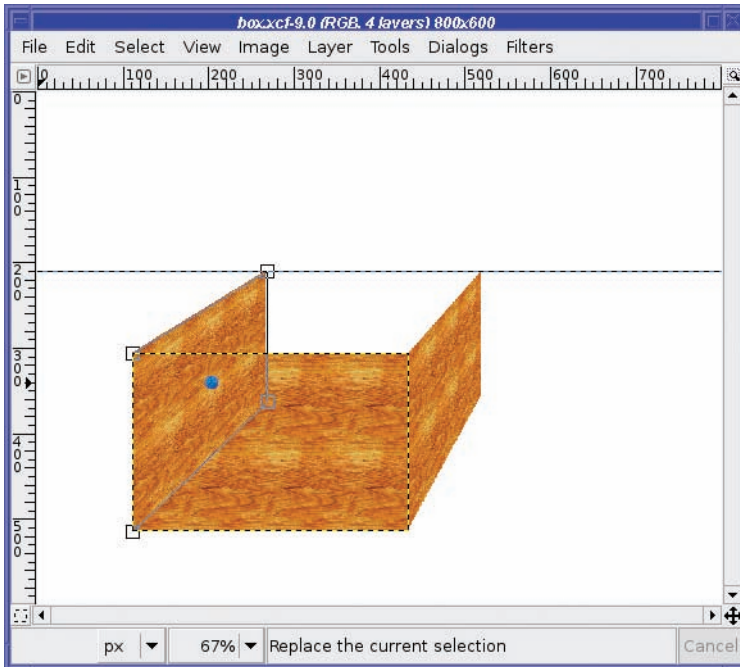
To get a horizontal guide, click on the ruler at the top of an image window, and drag downward (Figure 4-37). For a vertical guide, click on the ruler at the left side of the window and drag to the right. When you're moving a guide, the GIMP will automatically select the Move tool, so you may have to switch back to the tool you were using previously.



**Figure 4-37.** A horizontal guide

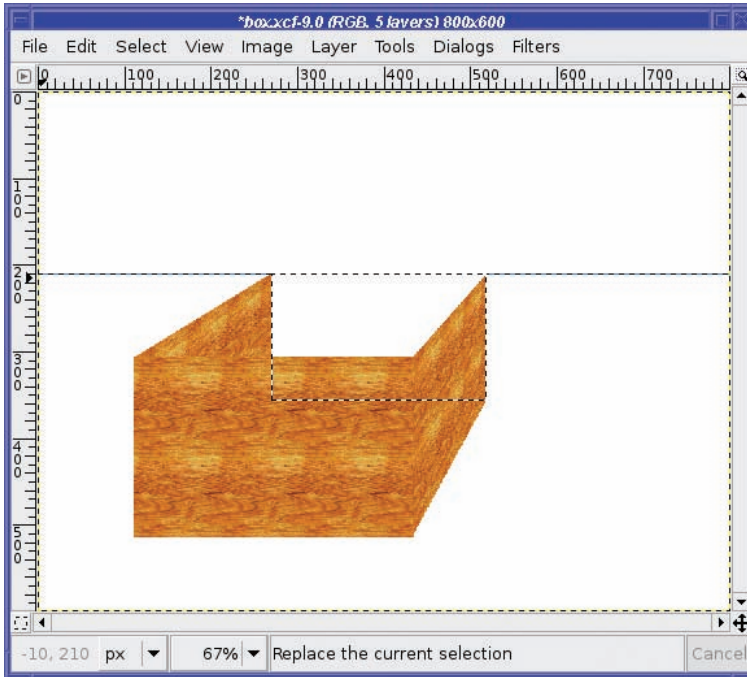
In addition to giving you a visual indication of where you are, guides have another useful property: when you drag near them, the mouse will “snap to the guide.” When moving a layer or working on a transform, this makes it easy to get the exact right spot. (You can disable *Snap to guides* in Preferences, but you probably won’t want to.)

After the perspective transformation is finished (Figure 4-38), you’ll want to move the layer down in the layer stack, using the up and down arrows in the Layers dialog. The left side should be *behind* the front side, so it needs to be below the front side in the layer stack.



**Figure 4-38.** *Transforming the left side*


Finally, it’s time to make the back side of the box. Just create a new layer, use the Rectangle Select tool to make a selection that goes from the back end of the box’s left side to the back end of the box’s right side (Figure 4-39), and then fill it. Move the layer down below the box’s front and right side layers. If you have trouble matching the new selection to the corners that are already there, it may help to set up vertical guides so that your selection will snap to them.



**Figure 4-39.** *Selecting the back side of the box. It's ready to be filled.*

It's not important whether the bottom of the selection matches anything: the bottom will be hidden by the front side of the box, so no one will ever see it. You could hide anything you want at the bottom of your box, and no one would ever know—except other GIMP users who look at your XCF.

You have all four sides of your box (Figure 4-40)! But it still doesn't look right, does it? The problem is lighting. A real box will have some sides brighter, others darker, depending on where the light is.

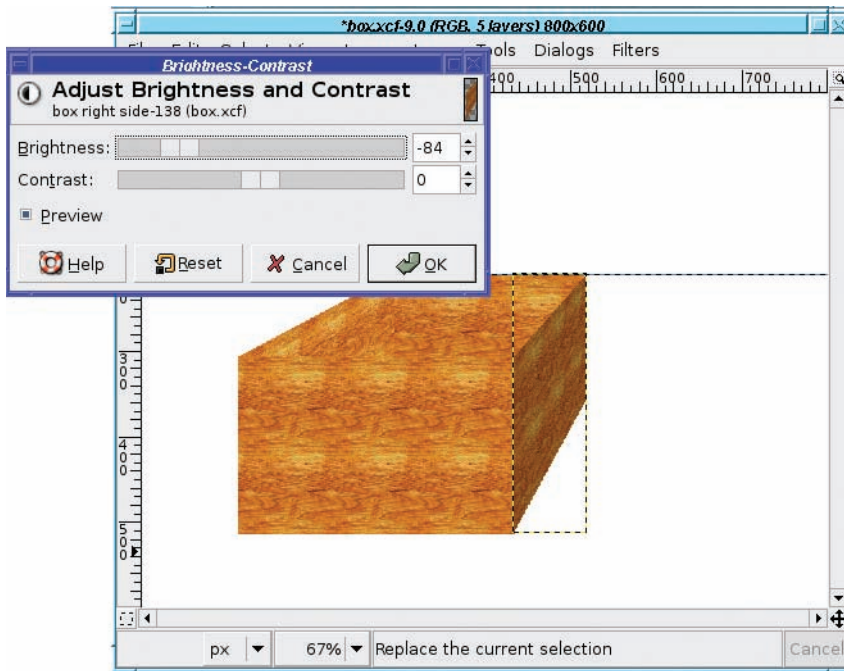


**Figure 4-40.** *All four sides of the box*

Your tree had the shadow on the right side of the trunk. If the box will be in the same image, it needs to match that. You need to leave the front of the box bright, but darken the left and right sides since they are both in the shade.

You already know how to make something darker using Brightness-Contrast, Levels, or Curves from Chapter 2. These operations work on one layer at a time, so it's easy to adjust each part of the box individually.


Select the layer for the box's right side, and then bring up your favorite brightness tool. (Which one you choose isn't important. You won't be making any subtle changes here, just general brightness.) I'll use the simplest one, *Layer > Colors > Brightness-Contrast*. Then make that layer darker (Figure 4-41).



**Figure 4-41.** Making the right side darker

Now darken the left side as well. You don't have to darken it exactly the same amount as you did the right side: with a real box, the darkness of the two sides would depend on the color of the table it's sitting on and the color of whatever's inside it. (Really! Try it with a real box and prove it to yourself.)

You'll probably want to make the back a little darker than the front, too, if only to keep it from blending into the box front where the two layers meet. But play around with the four colors until it looks like a box to you (Figure 4-42).



**Figure 4-42.** *The box after shading*

## Preparing the Planter Box

Finally, it's time to combine the two images and plant the tree in its planter box.

If you drew the box the right size for the tree, good for you! If not, no problem: use *Scale* to make them the right size.

---

**Tip** Make sure you're viewing both images at the same magnification setting.

---

Should you paste the box into the tree image, or the tree into the box image?

In order for the tree to go *into* the box, it has to be *behind* the front of the box, but *in front of* the back of the box, so the box needs to stay as multiple layers. The tree, however, can be pasted all as one layer. So it's easier to copy the whole tree using *Edit* ► *Copy Visible* in the tree image, then paste it into the box image, than it is to paste all those box layers one by one into the tree image.

*Wait!* Before you use *Copy Visible*, turn off the white background of the tree image. You want to paste a tree with a transparent background. Now aren't you glad you made a new layer when you started the trunk, instead of drawing it on the white background? (In Chapter 5 you'll learn how to remove objects from their backgrounds, but it's a lot more work.)

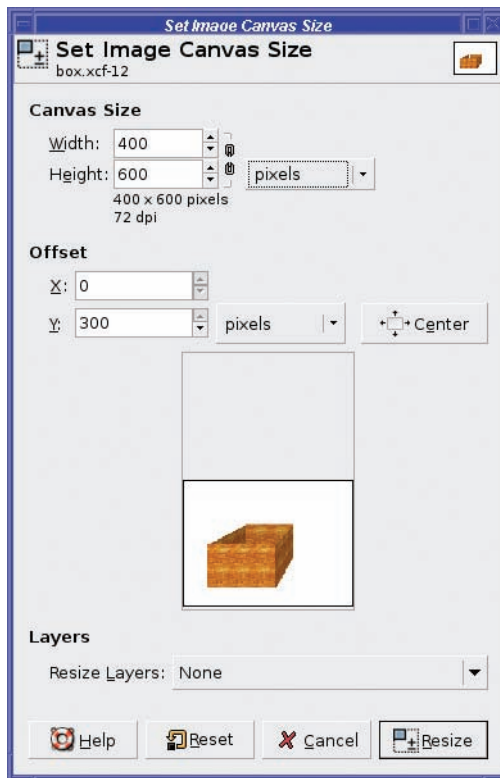
## Increasing Canvas Size

You've turned off the background and copied the tree. But there's one more problem to solve before you can paste: there isn't enough room in the box image for that tall tree!

Solving that is easy but it requires a few steps. First figure out how much space you need. You can use the Measure tool, or just combine the sizes of the two images. Too big is not a problem: you can always crop later.

Now resize the box image to be larger, using *Image* ► *Canvas Size...* This dialog's elements should look familiar to you now. It has a chain link icon, which you can click to unlink width and height if you want to increase just one or the other (as in this case). Go ahead and type in your new size now (or use the increase arrow for a *long* time). Then hit Tab or click in another text field to tell GIMP you're finished typing in the sizes.

Now you see something like Figure 4-43. There's a preview below the *Offset* fields, showing the new size of the image. It also shows the current location of the existing layers in the new, larger image. You can drag the image thumbnail around in this box.



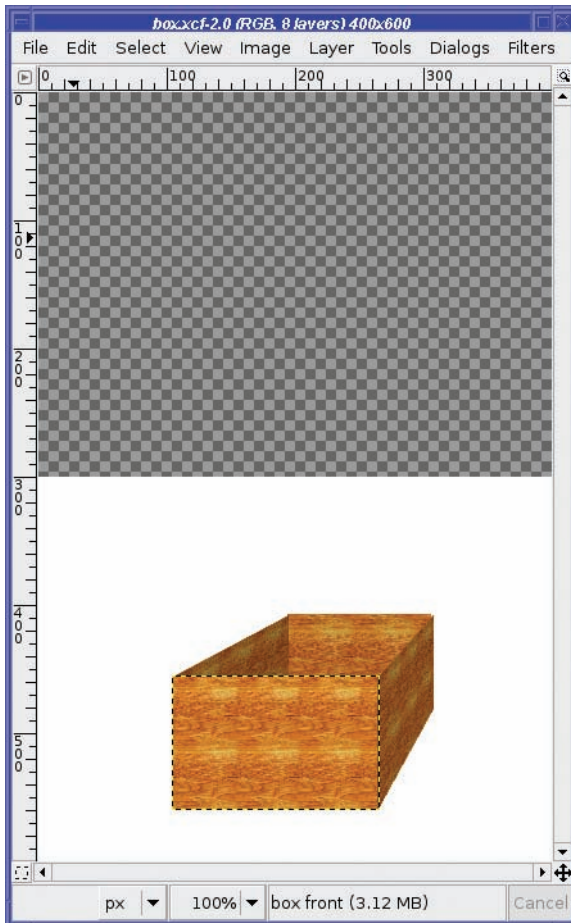
**Figure 4-43.** Changing the box image's size with *Canvas Size*

In this case, the goal is to get more space on top of the box, so drag the box thumbnail down to the bottom of the preview area. You can also type the offset of the layers directly in the *Offset* fields, but dragging the thumbnail around is usually easier.

If you're using GIMP 2.4, below the preview is *Resize Layers*, to specify which layers should be resized along with the image. Leave that at the default setting, *None*, for now; you'll see in a moment what it does.

Click *Resize* to get your new larger image (Figure 4-44).





**Figure 4-44.** Box image after increasing Canvas Size

The problem now is that the top-half of the image is blank. That's because the white background layer is still the smaller size.

To fix this, right-click on the background layer in the Layers dialog and choose *Layer to Image Size*. This resizes the background layer to fill the image.

If the background layer has no transparency, you're done: the layer will fill with the background color (white) and you're good to go. If the background layer has transparency, you'll need to *Fill with BG Color*.


Which brings us back to that *Resize Layers* menu in the Canvas Size dialog. (GIMP 2.2 and earlier didn't have *Resize Layers*.) When you adjust Canvas Size, GIMP can automatically resize *Image-sized layers* (all layers that are as large as the old image size), *All visible layers*, *All linked layers* (any layer where you've clicked the chain link icon in the Layers dialog), or *All Layers*. If you only need to resize one background layer, you may not need this option; but if you have multiple layers that need to be resized along with the image, this menu can be quite helpful.

## Plant the Tree

Finally, it's time to paste the tree into the box!

In your tree image with the background layer made invisible, *Edit* ► *Copy Visible* will copy all visible layers. Then paste into the box image, and click *New* to make it a new layer.

Put it *inside* the box: use the arrow buttons to move the layer down until it's above (in front of) the box's back and box's left side, but below (behind) the box's front and box's right side (Figure 4-45). Or just drag the layer's line to the right place in the stack.

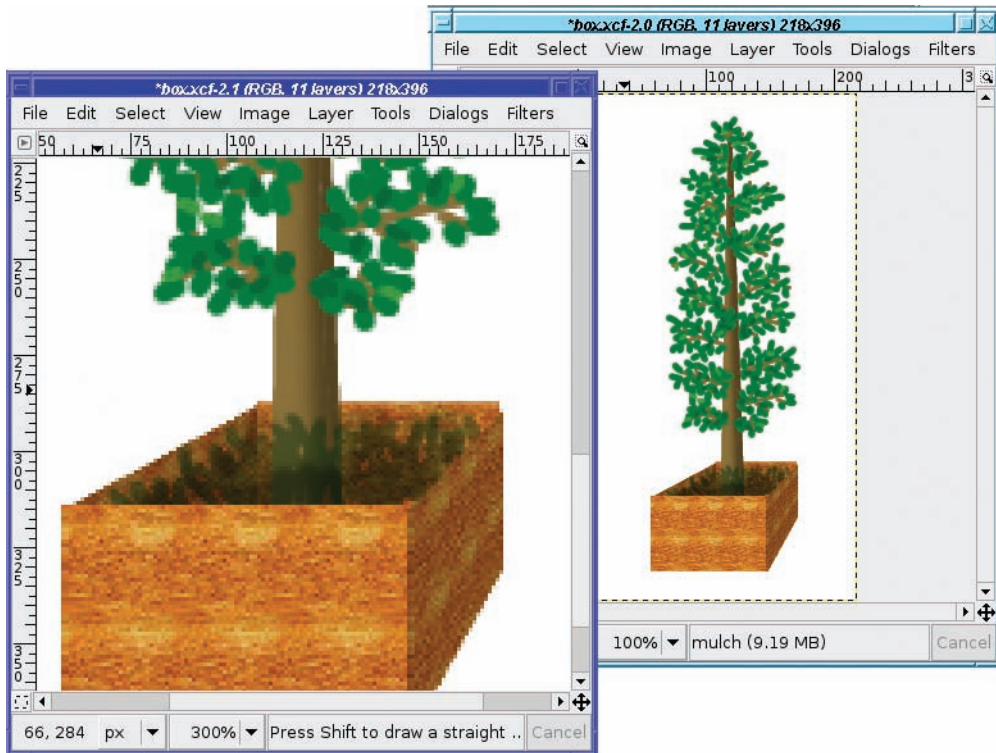


**Figure 4-45.** *Tree planted in box*

## Final Touch-Ups

You can stop here, if you like. But you can make the drawing look a little more realistic by adding some dirt or moss around the base of the tree.

Make a new layer just above the tree layer, with opacity at maybe 50% or so. Choose a dark green or brown color, and try scribbling some grass or mulch in the layer. You might also try adding some black (Figure 4-46).



**Figure 4-46.** *Adding some grass and mulch*

Finally, try adding a shadow. There are fancy ways of drawing shadows based on the shapes already in the image (you can use the Rotate and Perspective tools), but even adding a rough shadow drawn with the Free Select tool will make most images look more three-dimensional.

Start by making a new layer, named “shadow.” Put it above the background, but below all the other layers. Give the layer some transparency, maybe about 50-70%, and set the foreground color to black.

Now choose the Free Select tool and turn on feathering. Shadows always have fuzzy edges, but you’ll have to guess how fuzzy. Try to eyeball where the shadow should be (it may help to put a real physical box on the table in front of you and see where its shadow falls). Then use the Free Select tool to drag out the shape of the shadow.

Once you have a shadow shape selected, you can fill it with black (which will look fuzzy gray, not black, because of the layer transparency). But for a little more realism, try filling it using the Gradient tool. Make the shadow start out black close to the box, but make it lighter when farther away. Experiment with different shadow shapes, and with starting the gradient at different places to see the different types of shadows you can generate. Your tree is planted (Figure 4-47)! You're now an expert GIMP artist.



**Figure 4-47.** *The final tree in a planter*

## Summary

Now you know how to draw on layers, and have an idea when to make a new layer. You can draw lines, curves, and shapes such as rectangles and circles, and you can outline them or fill them. You've seen the various types of brushes GIMP offers, and how to use them with each of GIMP's painting tools. Finally, you've learned techniques of drawing, shading, filling with patterns, and using the Perspective Transform tool to model real-world objects.

You can use these techniques to draw just about anything you can think of. Just use a little patience, and a lot of layers.

Armed with this knowledge, it's time to go back to looking at photographs and learning how to select specific pieces so that you can use them in your own projects.



# Selection

**Y**ou've already worked with the GIMP's simplest selection tools: Rectangle, Ellipse, and Free (Lasso) Select.

So far, you've used selections primarily for drawing. But selections are more commonly used to choose one part of an image and ignore the rest.

This comes up in all kinds of applications. Perhaps you have a photo of a person against a complicated background, and you want to make the person stand out more—or you just want to paste a portrait onto a simpler background.

You may want to change part of an image—remember how the Bucket Fill tool could change the sky but not the rest of the photo, in Chapter 4?

Sometimes it's just a simple touch up. For instance, you might need to take black text on a white background and change it to green text on a yellow background. These are all selection issues.

If you've tried to use the Free Select tool to cut out parts of an existing image, you probably reached its limitations quickly. It takes a very steady hand, and if you make a mistake you have to start all over again.

Never fear! GIMP has quite a few tools for selecting shapes in existing images, in as much or as little detail as you need.

In this chapter, you'll learn about the following:

- Working with selections
- Select by Color and Fuzzy Select
- Bezier paths
- The Intelligent Scissors tool
- Modifying selections with selection modes
- The QuickMask
- Highlighting foreground objects
- Using channel masks to save a selection
- Layer masks
- Extracting foreground objects with SIOX

## Working with Selections

You know how to copy a selection and paste it somewhere else, but what else can you do with a selection?

### Marching Ants

When you have a selection in an image, be aware of the selection boundary. It's a black-and-white dotted line outlining whatever is currently selected. The segments along this line move, making the selection outline seem to shimmer. This changing outline is dubbed the *marching ants*, and you'll see the term used a lot.

The marching ants are helpful for telling you when you have a selection and where it is. But they can also get in the way, particularly when you're trying to blend in an object pasted from somewhere else, as you'll see. For this reason, it's very helpful to be able to hide the marching ants using *View* ► *ShowSelection*.

Squashing ants is used often enough that you may want to learn its keyboard shortcut, Ctrl+T (T stands for *toggle selection boundaries*). Another Ctrl+T will bring the ants back.

As useful as this is, though, it's also a major source of confusion. Once you've hidden the marching ants, you have no way of knowing whether anything is selected. If you try to conduct a GIMP operation that works only on the selection (which means *most* GIMP operations), you may be surprised when nothing happens.


---

**Caution** When you hide the selection boundary, try to remember to show it again as soon as possible. If you forget, you may regret.

---

### The *Select* Menu

Take a quick look through the *Select* menu (Figure 5-1), which contains a number of operations affecting the current selection.



**Figure 5-1.** *The Select menu*

*All* and *None* select the whole image, or make sure nothing is selected. You'll use these operations a lot, and it may be worth learning the keyboard shortcuts for them (Ctrl+A and Ctrl+Shift+A, respectively).

*Invert* reverses the selection: whatever is selected becomes deselected, and vice versa. You'll see why this is useful later in this chapter.

*Float* cuts out the selected area (leaving white or transparent behind) and pastes it as a new floating layer. Dragging inside a selection will do the same thing while also moving the selection. *Float* seems very similar to cut then paste. So why use it? One reason: *Float* keeps the selected area in the same location in the image. *Paste* usually centers whatever you're pasting, losing any information about where the selection was.

*By Color* will be discussed later in this chapter.

*From Path* converts the current path to a selection.

*Selection Editor* brings up a dialog with a few options that mostly duplicate other functions.


*Feather* makes the edges of the selection fuzzy.

*Sharpen* does the reverse, making a fuzzy selection sharper. However, if your selection is fuzzy because you've feathered it, you're best off using *Undo* to get back exactly what you had before.



*Shrink* and *Grow* make the selection larger or smaller.

*Border* replaces a selection with a new one that follows the border of the old selection (Figure 5-2). You can fill this with a color or pattern using the Bucket Fill tool, to create an effect much like stroking the selection; or you may find other uses for it.



**Figure 5-2.** *Bordering a selection*

*Rounded Rectangle* rounds the edges of a rectangular selection. Making a rectangle with feathering is another way of rounding a rectangular selection; but *Rounded Rectangle* gives you nice rounded edges that are still sharp, while feathering will make the edges fuzzy. (You could sharpen the edges afterwards, but that's more work.) You'll learn about the *QuickMask* and *Channels* later in this chapter.

*To Path* converts a selection to a path. Then you can edit it using the *Bezier Paths* tool (explained later in this chapter), stroke it, or save it in a vector-graphics format.

## Moving Selections

You may have already discovered, probably by accident, that dragging in a selection moves its contents, leaving a "hole" behind. If this ever happens unintentionally, GIMP has a fancy way

to cancel a mouse operation. Press the *right* mouse button while you're still dragging with the *left* button. Keep holding the right button down as you release the left. Whatever you were dragging should pop back to where it started. (You can release the right button now, too.)

Hitting Escape while dragging may also work in some versions of the GIMP.

Of course, if you forget these, undoing always works.

If you want to move the selection *boundary* without taking the contents along, use Alt-drag (as discussed in Chapter 3).

---


**Tip** If your computer uses Alt-drag for something else, Alt-Shift-drag will also move a selection boundary. Alt-drag may not work on Mac versions of GIMP. If it doesn't work for you, you can also move a selection boundary using the Move tool by setting its mode to *Transform Selection*.

---

## Select by Color and Fuzzy Select

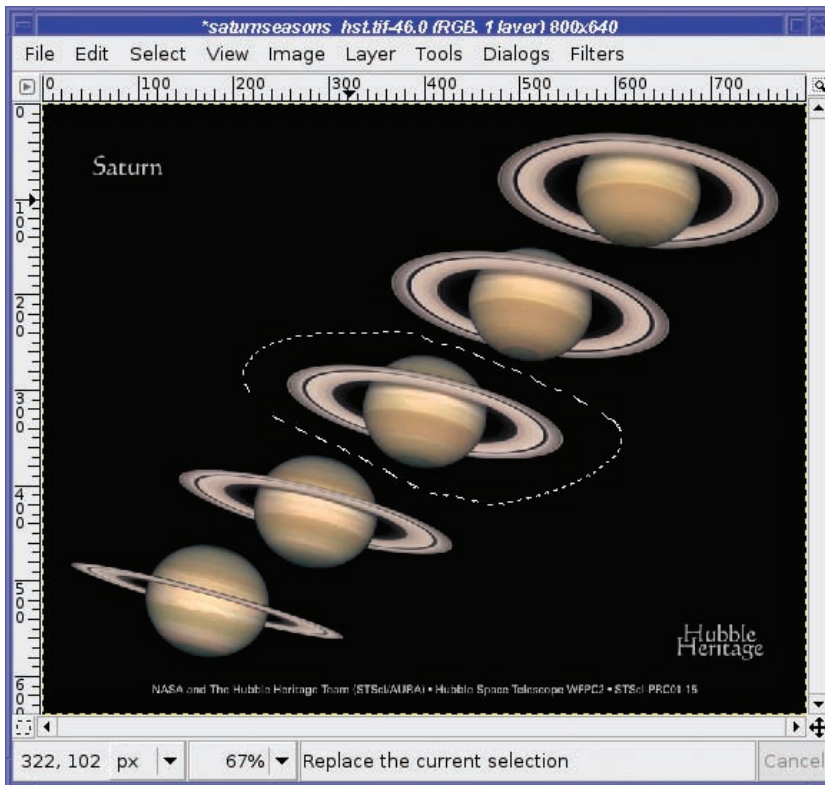
Select by Color is a very valuable tool in your selection arsenal. You can use it on images where you want to separate a foreground from a simple background.

For instance, the Hubble Space Telescope folks offer a poster showing several views of Saturn (Figure 5-3). I'd like to cut out the central image to make a Saturn that I can paste anywhere.



**Figure 5-3.** *Saturn, from the Hubble Space Telescope*

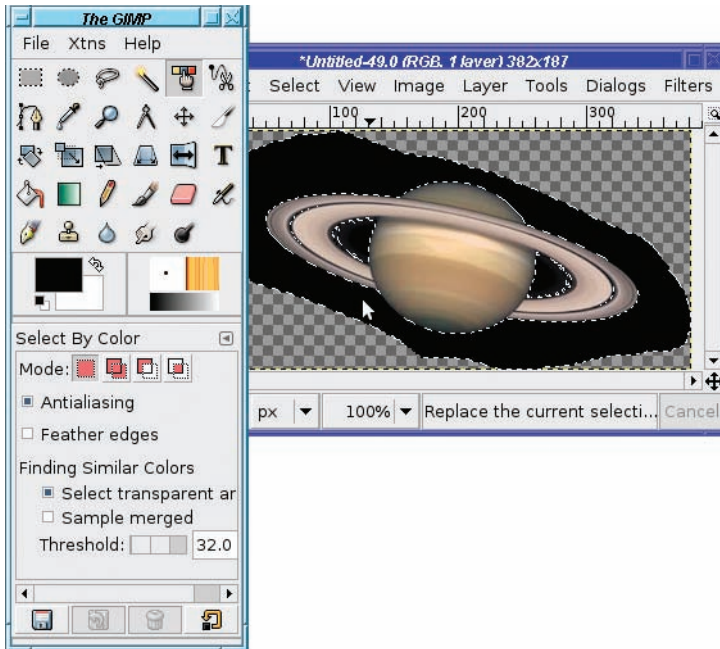
The first step is to single out just one of the Saturns. Since that doesn't involve any fine detail, I'll use the Lasso (Free Select) tool (Figure 5-4), and then use *Copy* and *Paste as New*.



**Figure 5-4.** Separate the middle Saturn from the rest with the Lasso tool.

Now it's time to use Select by Color. There's a button for it in the Toolbox (Figure 5-5), or you can use the *Select > By Color* menu.

Click in the image, on the color you want to select. In this case, I'll select the black sky: I'm trying to select everything that *isn't* Saturn. Why? Because that's easier than selecting Saturn. The sky is all one color; Saturn isn't.



**Figure 5-5.** *Select by Color*

Notice the tool options.

The *Finding Similar Colors* options work the same way as they did for Bucket Fill in Chapter 4. With *Threshold* at 0, only colors exactly the same as the one you click will be selected. With a higher threshold, GIMP will include colors that are similar but slightly different.

If you don't like the threshold when you first click, change it a little and click again. (If you undo first, the operation won't clutter up your undo stack, but it's not mandatory.) When you're trying to select an area that isn't all exactly the same color, like a sky, *where* you click in the image can make a big difference in what gets selected. Try clicking in different places with the same threshold, to see how that changes the selection.

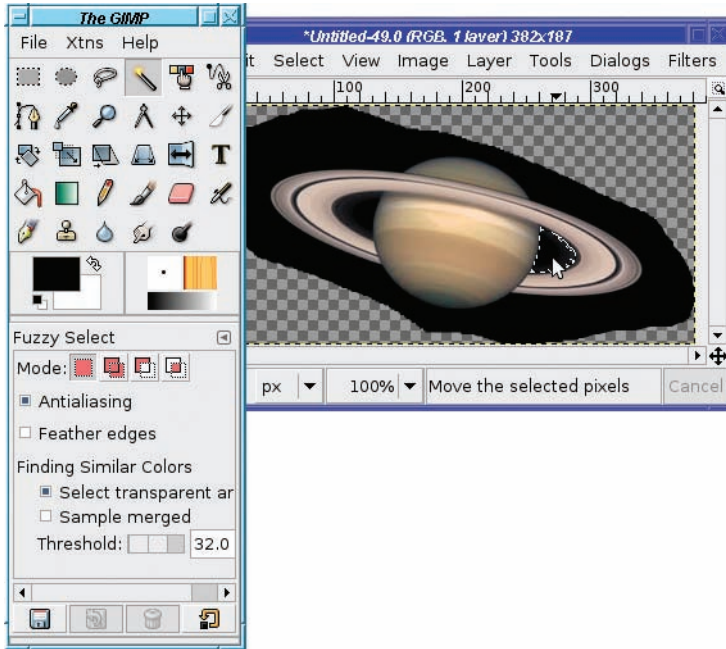
In addition, *Select by Color* offers the usual selection tool options you used in Chapter 4, *Antialiasing* and *Feather edges*.

Now that everything that *isn't* the planet is selected, you can make a transparent Saturn image by clearing with *Edit* ► *Clear* (Del in GIMP 2.4, or Control+K in earlier versions). This erases the black sky you just selected, leaving transparency instead (assuming the image has an alpha channel). Alternately, you can select Saturn and nothing else. Remember that *Select* ► *Invert* menu item, which selects everything that isn't selected, and vice versa? If you use it now, you can copy just Saturn. Either way, you end up with a Saturn that you can paste into other images.

Besides selecting solid-color backgrounds, *Select by Color* is particularly useful for extracting text or other single-colored elements out of logos, screenshots, or other computer-generated images.

## Select Contiguous Regions

Select Contiguous Regions, also known as Fuzzy Select or the Magic Wand tool (from its Toolbox icon) is very similar to Select by Color, and its tool options are basically the same. The main difference is that it only selects areas that touch each other (Figure 5-6).



**Figure 5-6.** Select Contiguous Regions

There are two other very important differences between Fuzzy Select and Select by Color. First, in Fuzzy Select, if you get the threshold wrong and want to try again, you have to cancel the previous selection (with Undo or *Select ► None*) to try again. It's all too easy to forget. If you don't undo, you'll end up moving the contents of the first selection you made.

The other difference is that with Fuzzy Select you can fine-tune the selection by dragging with the mouse. The marching ants appear as soon as you click the mouse down in the image. If too much or too little is selected, dragging down and right will increase the amount selected, while dragging up and left will decrease the amount.

Like Select by Color, Fuzzy Select is quite useful for text, especially when you need to select a single letter.

## Bezier Paths

The Bezier Paths tool lets you define *control points* (also called *anchor points* or *nodes*) outlining your object. Then it “connects the dots” to create a path. You can even bend the lines between the dots to create a smooth curve.



**Note** *Bezier* curves, pronounced *BEZ-ee-ay*, are named after Pierre Bezier, a French engineer who developed the concept while working on computer systems for automotive design.


A Bezier path isn't a selection itself. But once you have a path, it's easy to tell GIMP to select everything inside.

You can also use paths for other purposes. You can stroke a path, just as you can a selection. You can even export it to a vector-graphics format such as SVG.

However, for the purposes of this chapter, you'll mostly use paths to define the outlines of a selection.

## Defining a Path

The easiest way to understand Bezier paths is to see them in action. So open a photo and choose something you want to cut out of it. Then select the Bezier Paths tool in the Toolbox (Figure 5-7).



**Figure 5-7.** *The Bezier Paths tool*


I'm going to cut a single leaf from this photo so that I can paste it into other images. (It's poison oak! Picking a leaf and taking it home was *not* an option.)

Rule 1 for making a good path is: zoom in. Zoom *way* in. You probably want a zoom factor of 200% or more, depending on hand steadiness and how good your close vision is. Zooming in lets you place your control points *exactly* on the edges of the object you're choosing.

It doesn't matter if the object you're trying to select becomes larger than the window. The biggest advantage of Bezier paths versus a simpler tool like Lasso Select is that you can stop between points. That means you can rest for a while (making a complicated selection can be a tiring and time-consuming operation). You can also scroll the image or resize the window as much as you need to see exactly what you're selecting. Best of all, if you don't get it perfect the first time, you'll be able to go back and edit it later.

Time to get started!

Just click in the image window, selecting points you want to connect (Figure 5-8). Use as many or as few control points as you need, depending on how precise you want your selection to be.



**Figure 5-8.** *Beginning a Bezier path*

Keep clicking to place points around the outline of the image until you get back to the starting point. You can zoom and scroll as much as you need between points.

If you place a control point in the wrong place, you can click on the point again and drag it to the right place. You can even go back and click on points you added earlier, and drag them to new positions. As always, watch the cursor carefully as you move your mouse near the point.




When it changes to show crossed arrows, similar to the Move tool, that's your cue that you're over an existing control point and can drag it. (However, it can be difficult to tell the crossed arrows cursor from the simple cross cursor you see when you're adding a new point. If you look closely enough, you'll see the difference.) Undo (Control+Z) also works to remove the last point you added, and further undos will remove earlier points.

When you have gone all the way around, Control-click on the starting point (hold the Control key while clicking on it) to close the path.

You don't actually need to Control-click if you're just making a selection. Converting a path to a selection will usually work just fine if you just click normally on the original point. But for other operations such as stroking, it's a good idea to get in the habit of closing paths. It never hurts to Control-click on the first point once you get back to it.

## The Paths Dialog

In addition to being visible in the image, your path appears in the Paths dialog. That dialog is usually docked as a separate tab in the same window as the Layers dialog (Figure 5-9).



**Figure 5-9.** The completed path, showing the Paths dialog

Most of the components of the Paths dialog should be familiar to you from working with layers. You can manage multiple paths just as you would manage layers in the Layers dialog. Each path has the following:

- A small preview showing the general shape of the path
- A visibility icon (the “eyeball”)
- Space for a “chain link” between the eyeball and the preview, so you can link several paths and move them together

Most of the time, you probably won’t need to use more than one path at one time, but it’s good to know you can link them if the need should sneak up on you.

Paths are saved when you save to GIMP’s native file type, XCF. So if you spend a long time creating a path, it’s a good idea to save the image as XCF.

You can *Duplicate* a path (make a copy of it) in case you accidentally destroy what you spent so much work creating. Of course, there’s a *Delete* button for paths you no longer need.

*Selection to Path* converts an existing selection to a path. This is useful if you want to edit the path further, or use the shape for some other purpose, such as saving to a vector-graphics format like SVG. *Stroke Path* does much the same as *Stroke Selection*, which you used in Chapter 4.

But our immediate interest in the Paths dialog is *Path to Selection*.

## Creating a Selection from a Path

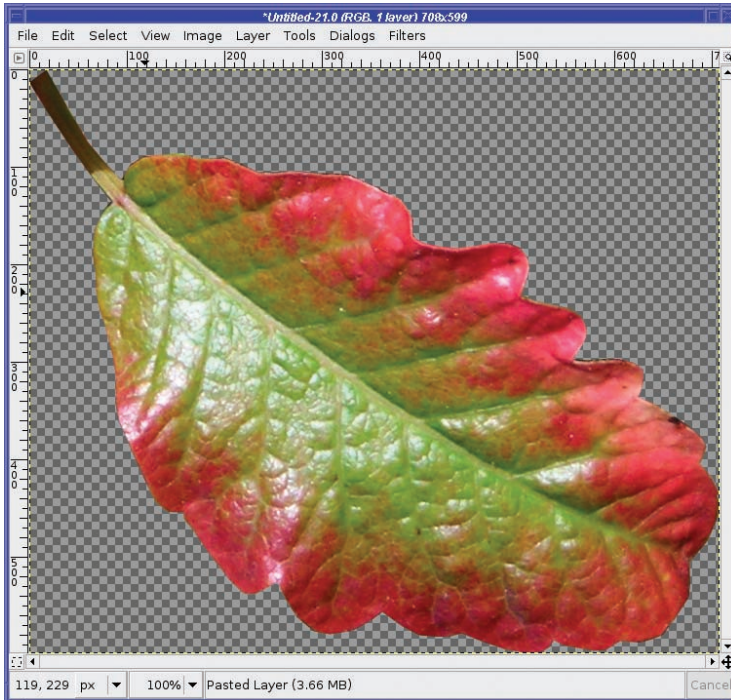
There are several ways to create a selection from a path.

The easiest way is to hit return. Voilà! You can also use the *Create Selection from Path* button in the Paths tool options, the *Path to Selection* button at the bottom of the Paths dialog (it has a pink square on it), the *Path to Selection* menu item in the context menu you get when right-clicking on a path in the dialog, or the *Select ▶ From Path* item in the image window’s menu.

Try it! Once you’ve defined your path, convert it to a selection. At this point, you’ll see both the path and the selection. If you find that distracting, you can hide the path by clicking the “visibility” eye button in the Paths dialog. Or simply switch to any other tool in the Toolbox (paths are only visible when the Paths tool is active).

To see what you’ve created, *Copy* the selection and (from the image window menus) open *Edit ▶ Paste as New*. GIMP will create a new image containing only your selection, on a transparent background (Figure 5-10).

Of course, you can also paste the selection into another image, or anywhere you want (Figure 5-11).




**Figure 5-10.** *The selection, pasted as a new image with a transparent background*



**Figure 5-11.** *The leaf pasted into another image*

The boundaries of the leaf are sharply defined, and it looks just fine pasted into other images. But often the boundaries don't work out quite so well, and the selection looks odd when pasted somewhere else. Other times, there is no sharp edge—imagine trying to select a photo of a shadow or a cloud.

You can also have problems if you make a relatively sloppy path. For instance, the path in Figure 5-12 doesn't follow the contours of the leaf very well.



**Figure 5-12.** A sloppy path

If you make that path into a selection and paste it into another image, the ragged edges show up and the result looks artificial (Figure 5-13).

You can make a sloppy selection look better by *feathering* it to make its edges softer. It's not a complete solution—but sometimes it's the best bet.

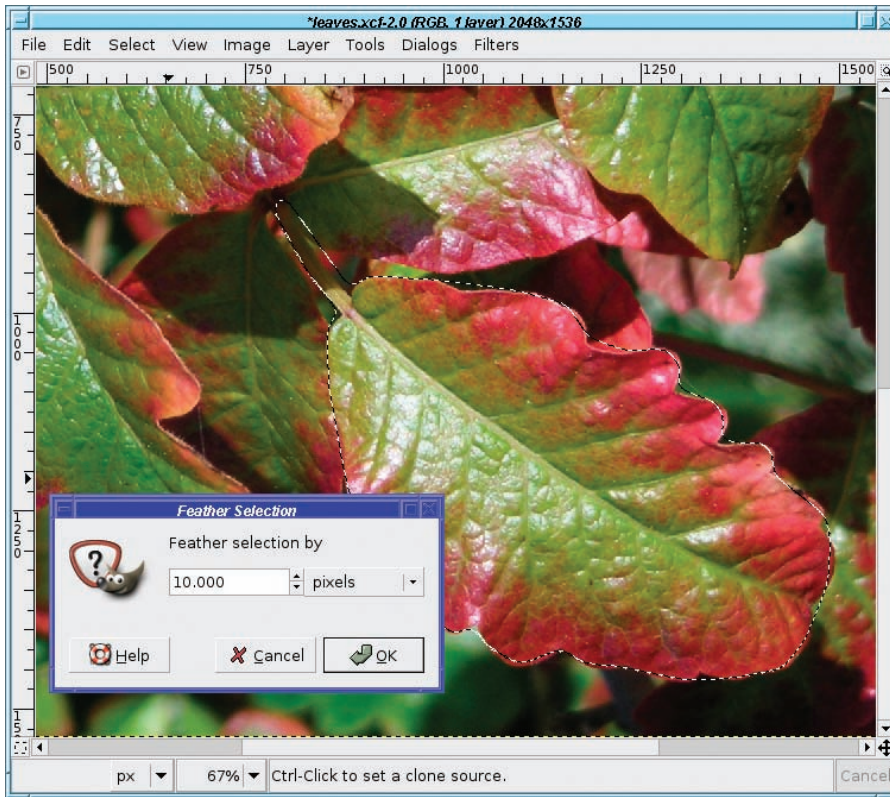




**Figure 5-13.** *A sloppy selection looks artificial when pasted into an image.*

Most of the selection tools (such as Rectangle, Ellipse, and Lasso) offer feathering as a tool option. But Paths doesn't, so you'll have to feather the selection after the fact, using *Select ► Feather...* from the image window. (Some people prefer feathering after selection anyway as it gives you more chance to increase or decrease the feathering as needed.)

This brings up the Feather dialog (Figure 5-14), which lets you choose the amount of feathering, specified in pixels.



**Figure 5-14.** Feathering a selection

The amount of feathering you need depends on how large your selection is, how ragged it is, and where you're pasting it. You will probably have to experiment and try several different settings to find the perfect value, undoing each feather before trying something different.

In this case, the default value of 10 pixels wasn't quite enough. A feathering of 15 pixels smoothed the selection enough for it to look okay when pasted (Figure 5-15). You can still see the problems around the edge of the leaf, but now they look more like shadows than image-editing artifacts.



**Figure 5-15.** *Feathering a sloppy selection makes it look better when pasted.*



## Curved Paths

So far, you've only used straight lines between path nodes. But a Bezier path can include curved lines as well. The key is the three edit modes in the tool options: *Design*, *Edit*, and *Move*. By default, you're in Design mode, which is best for adding nodes to a path or for moving existing nodes.

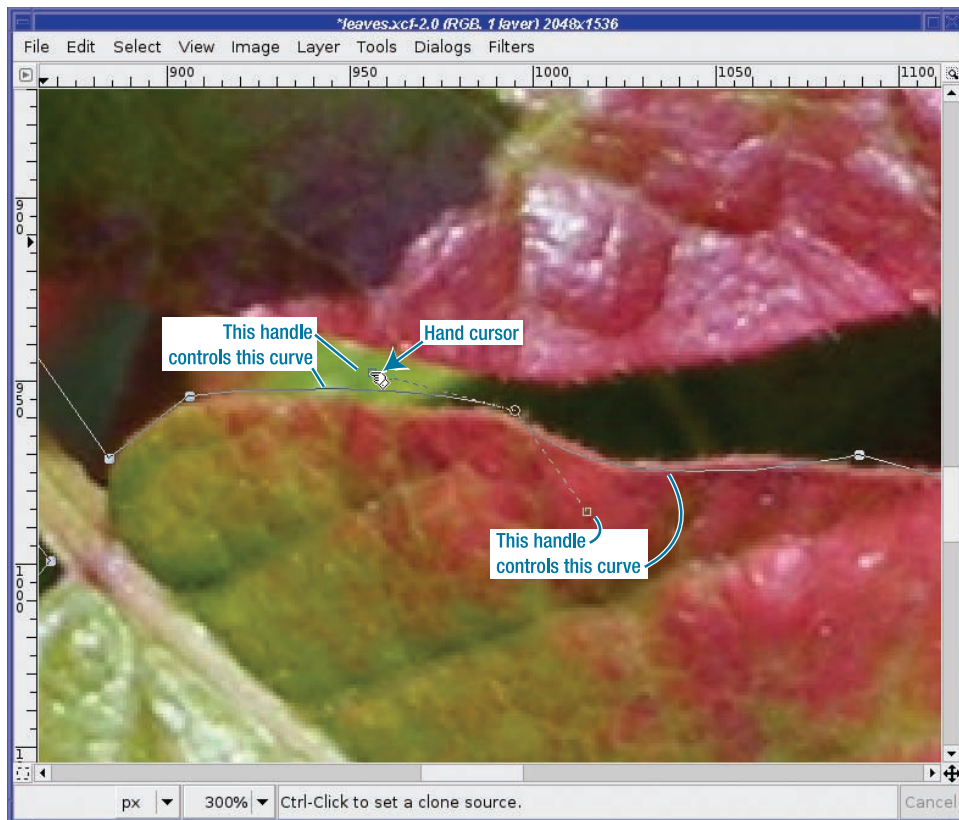
Each control point on a Bezier path can have two *handles* that control the curvature of the lines on either side. These handles are accessed via the *Edit* mode toggle button.

---

**Tip** Pressing the Control key while dragging on a control point will also switch to Edit mode and let you create handles for that point. As always, watch the mouse cursor when you press the Control key with the mouse over the control point. GIMP will tell you, via the mouse cursor, which mode you're in.

---

Figure 5-16 shows what the handles look like when you're using them. GIMP changes the cursor to a “pointing hand” when you're over a handle (though not all GIMP versions use the same cursors for editing paths, so watch what your own GIMP installation does). Drag the handle in the direction you want the curve to go. Experiment and get a feel for the sorts of curves you can make.



**Figure 5-16.** The handles for a Bezier control point

Basically, making the handle longer will make the curve stronger, while rotating the handle around its axis (the control point) will move the apex of the curve.

If you don't like the handles, you can just click on the path segment and drag it directly.

You can use paths that are composed of all straight lines, all curved lines, or any combination of the two. Usually, it's easiest to use straight lines for most of the path, and then curve only the few segments that follow tricky outlines. (Alternatively, you can just use more points for the tricky sections, and never use handles or curved segments at all.)

---

**Tip** A path consisting only of straight lines is called *polygonal*. You can prevent handles from appearing on new points by checking the *Polygonal* box in the Paths tool options. (If you have existing points whose handles you've adjusted, *Polygonal* won't remove their handles.)

---

## Adding Nodes or Segments and Moving Paths

Edit mode is useful for more than just curving line segments. You can also use it to add a new control point between two existing points. Just click along the path (make sure the cursor is a "+") wherever you want the new point.

This can be slightly confusing, because Edit mode assumes you want handles, so the new point will be created with two handles already visible. If you don't intend to use handles, you can avoid this confusion by checking the *Polygonal* box in the Paths tool options. Even if you *do* intend to use handles, you can still use this trick. Check *Polygonal* first, lay out your path, and then uncheck *Polygonal* and make handles for any control points you wish to use for curves.

You can also make a path that includes several disconnected path segments. Pay attention to which points are drawn with a closed circle versus an open or dotted circle. Any point that's open or dotted is *selected*.

Most of the time when you're creating paths, only one node, the last one created, will be selected. Creating a new point draws a line from the selected point to the new one.

But you can deselect a node by Shift-clicking on it, or select more than one node by Shift-clicking on each new point you want to select. If there are no nodes selected, or more than one, or a node that isn't the end of a path, then clicking in a new location will begin a new disconnected path segment.

Disconnected paths are useful mostly for making discontinuous selections or figures. They can be a time-saver since you can stroke or fill several areas at once. They're not generally used for making selections.

If you have multiple points selected, dragging any one of them will also move other selected nodes by the same amount.

## Moving or Modifying an Existing Path

The Paths tool has one more mode: Move mode. This allows you to move the entire path (not just one point, as in Design mode). Just click anywhere on the path and drag it to a new location. If you have more than one path segment, only the segment you drag will move. If you want all path segments to move, using Shift-drag will move them all at once.


If you change windows or choose another tool after working with a path, when you go back to the window the path is no longer visible. How do you go on working with an existing path?

First, make sure that the right path is selected in the Paths dialog, and that its visibility icon (the eyeball) is on. This will show the path, but not its control points, in the image window.

Now, with the Paths tool selected in the Toolbox, click on the path in the image. The control points should appear, and you're ready to continue editing it.

## The Intelligent Scissors

The Intelligent Scissors tool (Figure 5-17) works somewhat like Bezier paths, but with a twist: it tries to follow the boundaries of an object. Use it the same way you would the Paths tool: click close to where you want your selection boundary to be, and in theory, GIMP will make a more complicated path which follows the edges it can detect in the image. The result will be a selection, not a path.



**Figure 5-17.** *The Intelligent Scissors tool*

When you've finished defining the outline of your selection, click on the first point to close the path, and then click inside the path to select.

When it works, Intelligent Scissors can make a selection that's more accurate and detailed than you would likely make yourself using the Paths tool. But unfortunately, it doesn't always get it right. Do try the tool, but most people find the regular Paths tool more useful. If you find that the cursor sometimes “jumps” to put a point somewhere where you didn't want it to be, Shift-clicking might help that.

If the Intelligent Scissors tool gets close, don't give up! You can go back and move existing points or add control points to fine-tune your work, at least until you click inside the selection to “finish” it. Even after you finish the selection, you can still rework it by converting it to a path, or through other means such as the QuickMask (which will be discussed later in this chapter).

Intelligent Scissors only makes selections. You can't use it to make paths for other uses. Therefore, it has the normal selection tool options: *Mode*, *Antialiasing*, and *Feather*.

It also has an additional option: *Show interactive boundaries*. With this enabled, you can see as you mouse down where the tool would put the boundaries, then drag the mouse left, right, up, and down before you release the button to get the selection exactly where you want it. On anything except a very slow machine, you'll probably find this option helps a great deal.

## Modifying Selections with Selection Modes

It often happens that you use one of the selection tools and it gets it *almost* right—but there's a little bit that's wrong. If you could just add a little to the selection, or subtract a little, it would be perfect.

For example, I used Fuzzy Select on the sky in this image of the observatory domes on Mauna Kea (Figure 5-18).

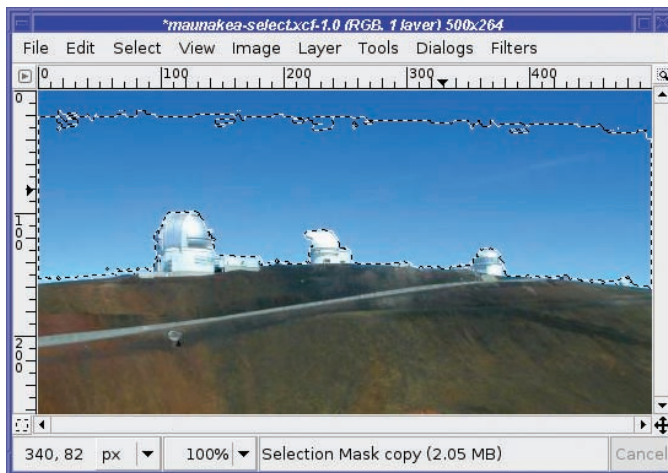
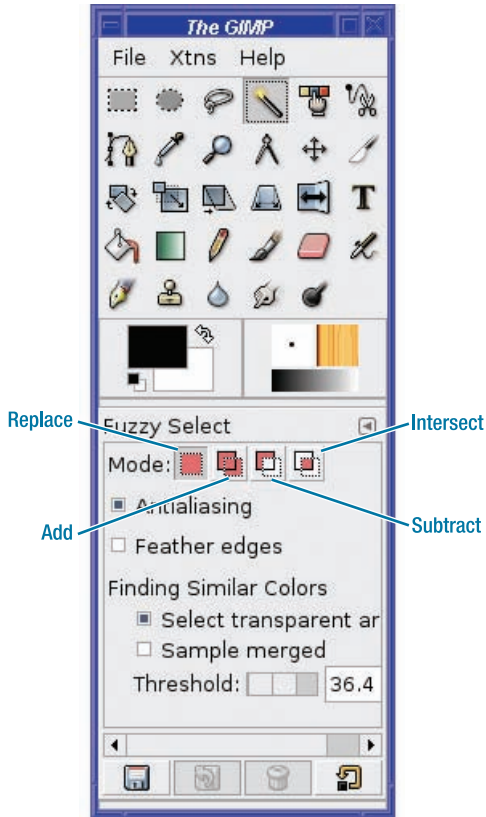


Figure 5-18. Sky partially selected with Fuzzy Select

The hard part—the selection around the domes and near the ground—is fine. But the very top of the sky isn't included. (Changing the Fuzzy Select's *Threshold* doesn't help—it just selects the domes and the road in addition to the sky.) If I could just add the top part of the sky to the existing selection, everything would be perfect.

*Modes* come to the rescue. Every selection tool has four modes (Figure 5-19) represented by buttons at the top of the tool options area.




**Figure 5-19.** Selection modes

The default mode, on the far left, is *Replace*. Making a selection in *Replace* mode removes the current selection, and replaces it with the new one.

## STARTING SELECTIONS

You may remember that with most selection tools, clicking *inside* an existing selection will move the selection contents, rather than beginning a new selection. So Replace mode only replaces if you start the new selection *outside* the boundaries of the old one. The same holds for the other selection modes: any new selection must usually be started outside the boundaries of the existing one. An exception is the Rectangle Select tool in GIMP 2.4, which does allow you to start a selection inside an existing one.

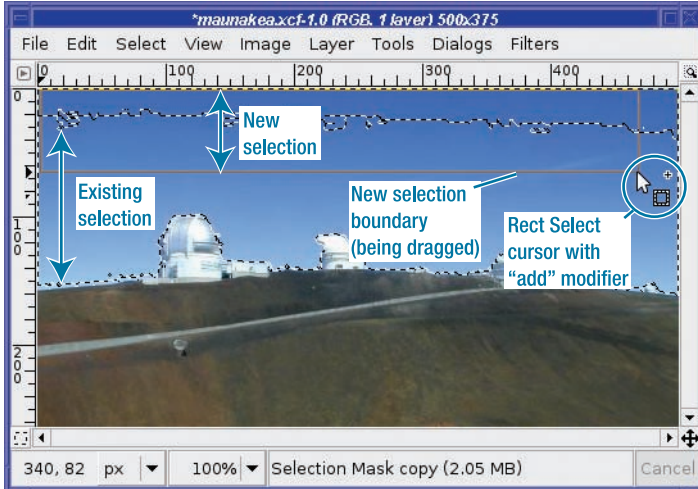


Watch the cursor carefully whenever you start a selection, if there's already a selection active. The "Move" component of the cursor will tell you when you're still inside a selection. If you're outside, the cursor will show the icon for the selection tool you're using (Figure 5-20). You can't depend on the marching ants to tell you: if the selection is feathered, it may extend beyond the marching-ants boundary.

*Add* mode adds the new selection to the existing one. In the Mauna Kea example, that's what I want: to keep the existing selection, but to add the sky at the top of the frame.

The easiest way to do this is to choose the Rectangle Select tool, and set the mode to Add. Then make a rectangular selection that includes everything from the top of the frame down into the existing selection (Figure 5-20). A fast way to do this in GIMP 2.4 is to start inside the existing selection and drag upward and over to one corner, and then adjust the other side of the rectangle so that it covers the other corner. With earlier GIMP versions, you can't do that because starting a new selection inside an existing one, even in Add mode, moves the contents of the existing selection; and you can't adjust the size of a rectangular selection once it's made.

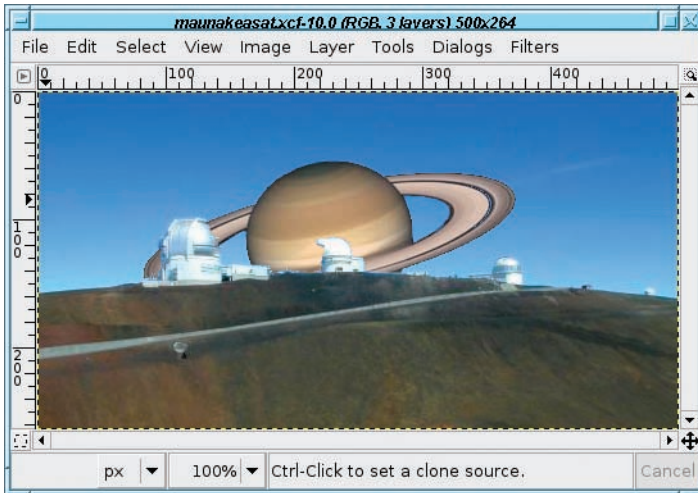




**Figure 5-20.** Use the Rectangle Select tool to add a selection at the top of the sky to the existing selection.

The two selections are added together. Consequently, the entire sky is selected.

Now you can do operations such as *Clear* to remove the sky. Or you can make the sky a separate layer (with *Float*, or just cut and paste it), put it at the bottom of the layer stack, and interpose another object between the foreground and the sky (Figure 5-21).



**Figure 5-21.** Interposing another object between foreground and sky



One way to obtain this effect is to invert the selection (*Select ► Invert*) so you'll now have the foreground selected. Copy, then paste the foreground and make it a new layer. After that, you can paste the transparent Saturn (or anything else you want) between the two layers.

Subtract mode is the inverse of Add mode: it subtracts from the selection. This is particularly handy when confronted with the opposite of your previous problem: an automated tool has selected just a little bit too much.

## USING MODIFIER KEYS FOR SELECTION MODES

Pressing the Shift key before dragging out a selection will enable Add mode. However, the Shift key also has another meaning in some selection tools: in Rectangle and Ellipse Select, it constrains the selection to be a square or circle. (If you find this confusing, use the Mode buttons in tool options when you want to add to or subtract from the selection.)

Then how do you constrain a selection to a square or circle if you *don't* want Add mode? The trick is to start dragging, and *then* press Shift. This will turn on the constraint without changing the mode.

Likewise, pressing the Control key will enable Subtract mode, but in some tools it has another effect. In Rectangle and Ellipse Select, Control makes the selection expand outward from the starting point, instead of using that point as the upper-left corner. The trick is to hold down the Shift key while you start the drag, and then release it while still dragging.

Intersect mode adds to the selection *except* where the old and new selections overlap: there, it will deselect what was selected previously.

## The QuickMask

Sometimes, outlining just isn't the best way to select an object. You may catch yourself wondering, "Isn't there a more direct way to say that I want to select *here* and *here* but not *there*?"


GIMP has an answer: the QuickMask.

The QuickMask lets you see the selection visually superimposed on the image instead of the usual "marching ants" view. Everything selected will look normal; everything *not* selected will have a translucent red mask over it.

You can paint on the mask with all the usual tools, like the Pencil and Paintbrush. Everywhere you paint in white will be selected; paint black to deselect. Painting in gray, or with the edge of a fuzzy brush, will partially select.

The QuickMask is one of those mind-bogglingly useful features that lots of people don't know about. It's easy to miss. I used GIMP for years before I stumbled onto it. Don't make that mistake, because the QuickMask can make complex selections easy. Just remember that red means *not* selected.

Flip to QuickMask mode by pressing the small square button at the bottom-left corner of the image window (Figure 5-22). You can flip back and forth between QuickMask mode and normal (marching ants) mode at any time with that button.



**Figure 5-22.** The QuickMask button. A rough selection of the biplane has already been made with Free Select.

---

**Note** Areas that are partially selected (such as the edges of a feathered selection) will have some red, but the mask will be more transparent than usual. This is not always easy to see, particularly on color images which themselves have red; but it still gives you a better idea than the marching ants would. You can make the QuickMask use a different color by right-clicking on the QuickMask button and choosing *Configure Color and Opacity*.

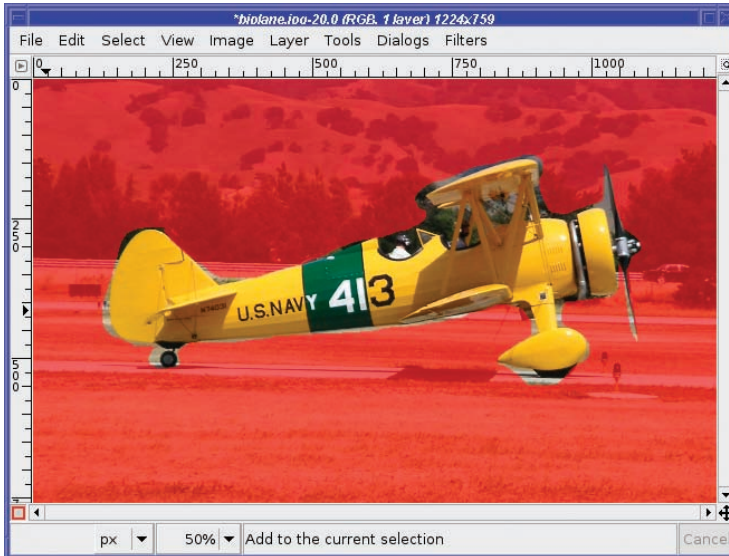
---

So how do you use the QuickMask?

It's often best to use some other tool to make an approximate selection, and then use the QuickMask to clean it up. For instance, I'll select the biplane, which you may remember from the animation back in Chapter 3. I'll begin by making a quick outline with the Free Select tool, as Figure 5-22 shows. My selection isn't very accurate; I just want something approximate.

Now press the QuickMask button. The marching ants disappear, and everything that wasn't selected before is overlain with a translucent red mask (Figure 5-23). Notice that the appearance of the button has changed, to show you that you're in QuickMask mode. Press the button again, and you'll go back to seeing the marching ants.


When it's time to draw on the QuickMask, Rule 1 is the same as for paths: *Zoom in!* You want to be able to see individual pixels, so you can paint precisely even if your hand isn't that steady.



**Figure 5-23.** *The QuickMask view of the biplane Lasso selection*

If you want a feathered selection (which is usually best for selecting an object out of a photograph), use a fuzzy brush. For a sharp-edged selection, stick to a sharp-edged brush. The size of the brush you choose will depend on the size of your image and the complexity of your subject.

Then start painting! (Figure 5-24.) Shift-clicking to draw straight lines will come in handy here.




**Figure 5-24.** *Painting on the QuickMask with the Paintbrush tool*

Use short strokes. That way, if you draw somewhere you didn't intend to, Undo will correct the mistake without losing much.

Remember, when you paint on the QuickMask, use black for areas you *don't* want selected. For areas you *do* want selected, use white or the eraser (assuming your background color is white). Using gray, or any other color, will only darken the QuickMask partially: the area will be partially selected.

You might think you can't select areas and fill them—after all, the QuickMask defines an existing selection! If you make a new selection, won't that destroy the one you're trying to refine?

Fortunately, GIMP lets you make “selections” and fill them while you're painting in QuickMask mode. So go ahead and use the drawing techniques you learned in Chapter 4, including making shapes to be filled with black or white (Figure 5-25).



**Figure 5-25.** Use the QuickMask to fill in gaps.

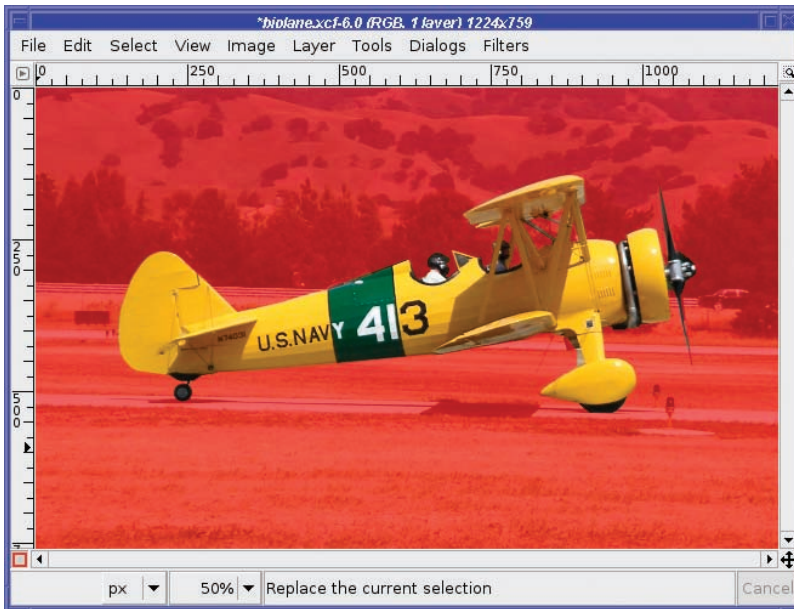
---

**Tip** It often helps to use a drawing tool such as the Paintbrush to outline an area, and then use Free Select to fill the remaining real estate. That way, you can fill large areas quickly using only crude Free Select strokes. Figure 5-26 shows a free selection about to be filled.

---


Continue painting, switching tools as needed to cover areas and fill gaps until you're happy with the result (Figure 5-26). Then click on the QuickMask button to return to the normal view.





**Figure 5-26.** The QuickMask is finished, and ready to be turned back into a selection by pressing the QuickMask button again.

Then you can Copy and Paste as New (Figure 5-27), or use your selection any way you choose.




**Figure 5-27.** The object selected with the QuickMask is now pasted as a new image.

## Highlighting Foreground Objects

What can you do with a selection, besides copying it and pasting it somewhere?

Pretty much anything, as it turns out. Nearly all of the GIMP's operations will operate only on the selection, if one exists. For instance, GIMP's brightness tools, discussed in Chapter 2, will lighten or darken only the selected part of an image.

Figure 5-28 has a background that distracts from the flower in the foreground. Can GIMP de-emphasize the clutter to make the subject stand out more?




**Figure 5-28.** *Nice flower, distracting background*

First, select the flower using any combination of the techniques already discussed. Select by Color gets most of it, and correcting with the QuickMask fills in the rest. Then invert the selection: since the goal is to change the background, you want the background selected, not the flower.

Now bring up your favorite brightness tool. In this case, Curves is particularly helpful, because it lets you reduce the brightness of the white wall with respect to the green leaves (Figure 5-29). After darkening the background, you may also want to try blurring it a bit with *Gaussian Blur* (see the section in Chapter 6 on Gaussian Blur for more details).



**Tip** You'll probably want to make the selection invisible with *View* ► *Show Selection*. Then you can make sure the boundary between the flower and the background isn't too obvious. If you can see a line at the selection boundary, or the object looks like it has a "halo" around it, you probably need to feather the selection more.



**Figure 5-29.** Use Curves to dim the bright background.

## Using Channel Masks to Save a Selection


Sometimes it would be really nice to be able to save a selection.

Perhaps you have a selection you've spent a lot of time refining, but you temporarily need to make another selection somewhere else in the image. Or maybe you've made a selection in one image, but you want to be able to use that selection in another related image.

Of course, you can convert the selection to a path, and give the path a name. But a path won't store all the information about your selection: if you've spent oodles of time feathering the edges just right, converting to a path will forfeit all that work.

The answer is a *channel mask*.

Channel masks are black-and-white images (similar to the QuickMask) that represent some aspect of an image. To work with channel masks, you need the Channels dialog, usually docked as a tab between Layers and Paths (Figure 5-30).



**Figure 5-30.** The Channels dialog

Notice that there are three channels already there, for the three colors in an RGB image: Red, Green, and Blue. If your image has transparency, there will also be a channel for Alpha. Each channel has a visibility button next to it, just like layers do. Click on the eyeballs to turn off individual colors and watch what happens to the image.

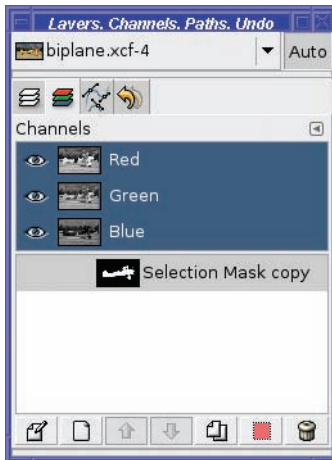
But you won't be using color channels now (for more on color manipulation, see Chapter 8). Instead, you'll be creating a special channel to hold the selection.

---

**Note** If you view the Channels dialog when you have the QuickMask active, you'll see that the QuickMask is also a channel.

---

To save a selection as a channel, use *Edit* ► *Save to Channel* (Figure 5-31).



**Figure 5-31.** The selection is now saved as a channel.

Notice that the selection channel doesn't have its visibility button turned on—by default it's invisible. You can show it by clicking in the spot where the “eye” would be.

The channel will appear as a translucent overlay, like the QuickMask, except it will be dark gray, not red. (You can change the color with *Channel Attributes*, as you'll see.)

Like the QuickMask, you can paint on a channel using the normal drawing and filling tools, with a few caveats:


- A channel must be visible (have an “eye”) in order to paint on it.
- If you still have a selection active, you won't be able to paint outside the selection. If you convert a selection to a channel, and then paint on it, you'll want to do *Select ▶ None*.

You can save as many selection channels as you wish. You can rename channels in the same way that you rename layers: double-click on its name in the channel dialog and type in a new one. Once you've saved a selection, you can make another selection and save that, too (Figure 5-32). The small preview in the Channels dialog will give you an idea of what each selection channel looks like.



**Figure 5-32.** Two selections saved (and renamed)

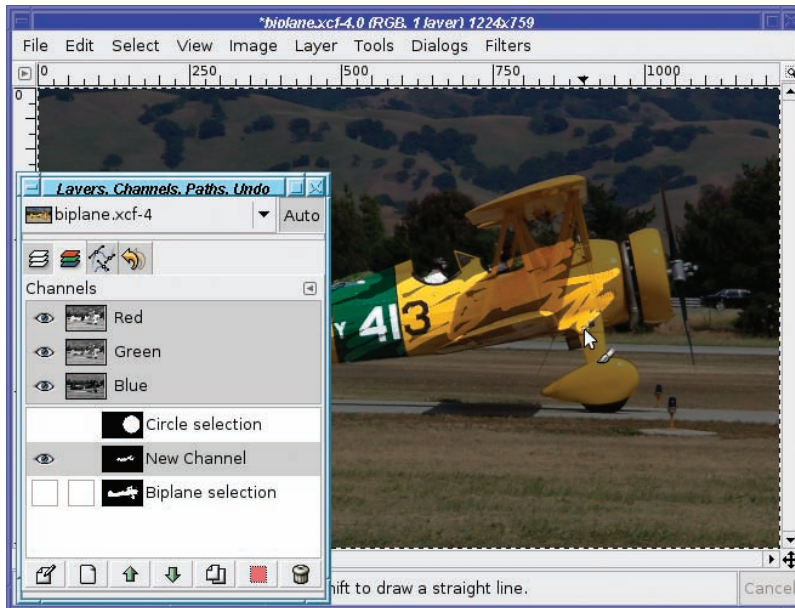
To convert a saved selection channel back to an active selection, right-click on the channel to get a context menu (Figure 5-33).



**Figure 5-33.** The Channel dialog's context menu

*Channel to Selection* is the simplest operation: just replace any existing selection with the selection saved in the channel. But you can also use the other selection modes: *Add to Selection*, *Subtract from Selection*, and *Intersect with Selection*. They let you combine several selections in fairly complex ways. Of course, you can also make a copy (*Duplicate Channel*) or *Delete Channel*.

You can also make a *New Channel*. By default, it will be all black (which will appear as a gray overlay), and then you can paint on it. In Figure 5-34, I've scribbled on the biplane's fuselage with the Paintbrush tool and a large brush.



**Figure 5-34.** *Painting on a new channel*

The first item in the context menu, *Edit Channel Attributes*, gives you another place to change the channel's name. In addition, it lets you change the *Fill Opacity* (Figure 5-35).

On some images, it may be nearly impossible to see what you're doing when the image is covered by the translucent dark gray of a channel. Changing *Fill Opacity* can make the overlay lighter or darker.

You can also change the color by clicking on the color box. Figure 5-36 shows the result.

Channels don't get in your way because you can switch them off at any time. They won't disturb your current selection until you explicitly ask for that using the context menu.

Of course, channels are included when you save as XCF, so all your stored selections will still be there the next time you edit the image.



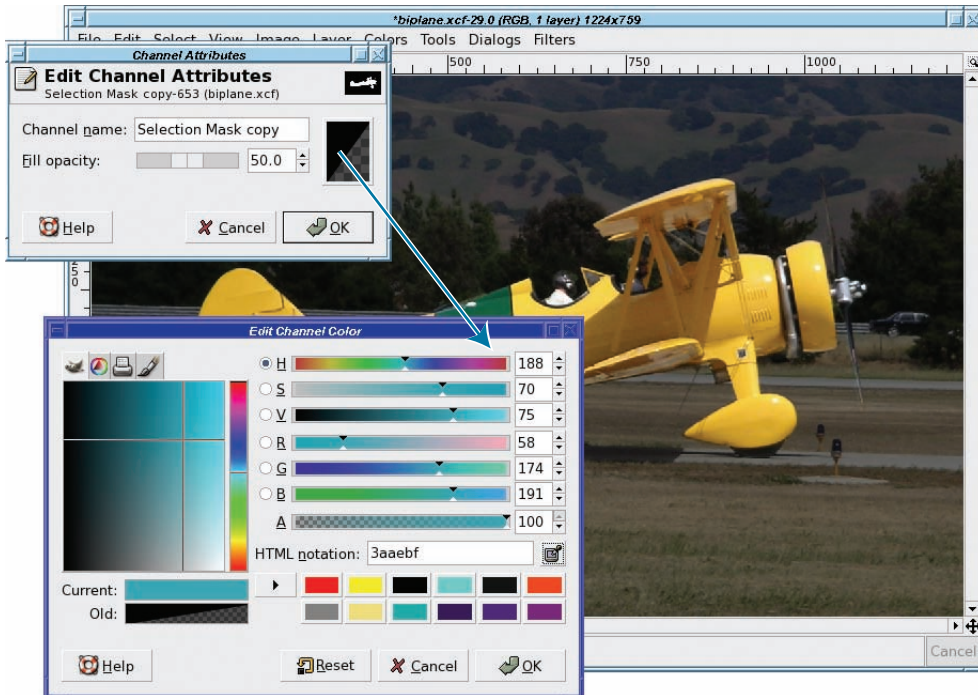


Figure 5-35. Channel Attributes. Click on the color button to bring up a color dialog.

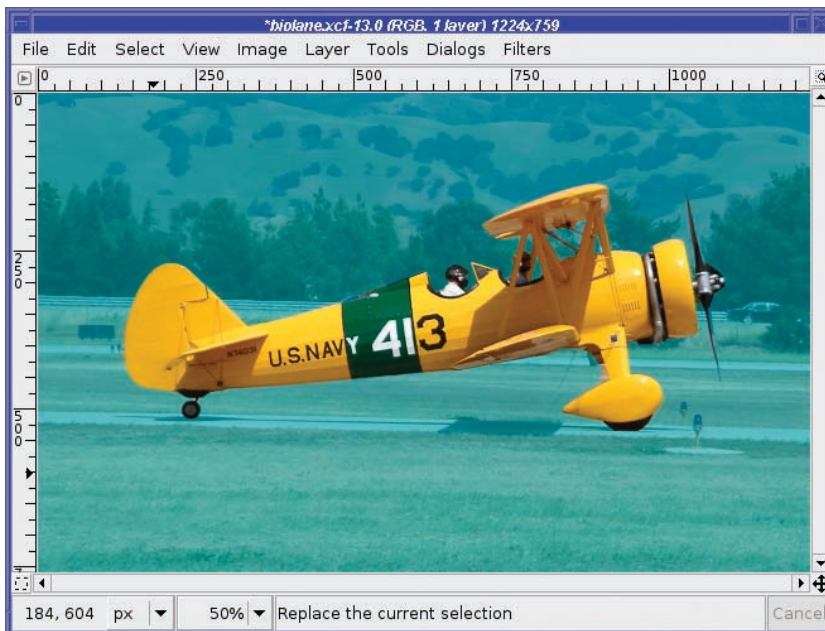


Figure 5-36. The channel's overlay is now a different color.



## Layer Masks

A layer mask defines how much, and which parts, of a layer will be visible.

Layer masks are not actually a selection tool. But you can use them in many of the same ways you'd use a selection. For some jobs, a layer mask may be easier and more intuitive than using a selection.

A layer mask is actually a black-and-white image, like a channel or the QuickMask. It represents the “alpha channel” of the layer. The whiter the mask is, the more the affected layer shows through—the less transparent it is. Wherever the channel mask is dark, the image will appear transparent.

It may help to think of the mask as a film negative placed on top of your image. Anywhere the film is black, you can't see your image; anywhere the film is clear (or in this case, white), the layer can be seen.

Layer masks are particularly useful when combining images. For instance, I have an ostrich (Figure 5-37) and a giraffe (Figure 5-38). What if I want to make an ostraffe?





Figure 5-37. An ostrich




**Figure 5-38.** A giraffe

Start by pasting the ostrich onto the giraffe image. Copy from the ostrich photo, paste over the giraffe, then go to the Layers dialog and click the *New Layer* button. Rename the layers so you'll remember which is which.

The next step is to line up the two images. But that's hard since you can't see the giraffe once the ostrich is on top of it. The *Opacity* slider in the Layers dialog saves the day. Set opacity to about 50% (or whatever makes it easy to see both images), and it's no problem to line them up (Figure 5-39).





**Figure 5-39.** Make the top image translucent in order to line up the two images.

Now it's time for the layer mask. Right-click on the ostrich layer in the layers dialog and select *Add Layer Mask...* (Figure 5-40).

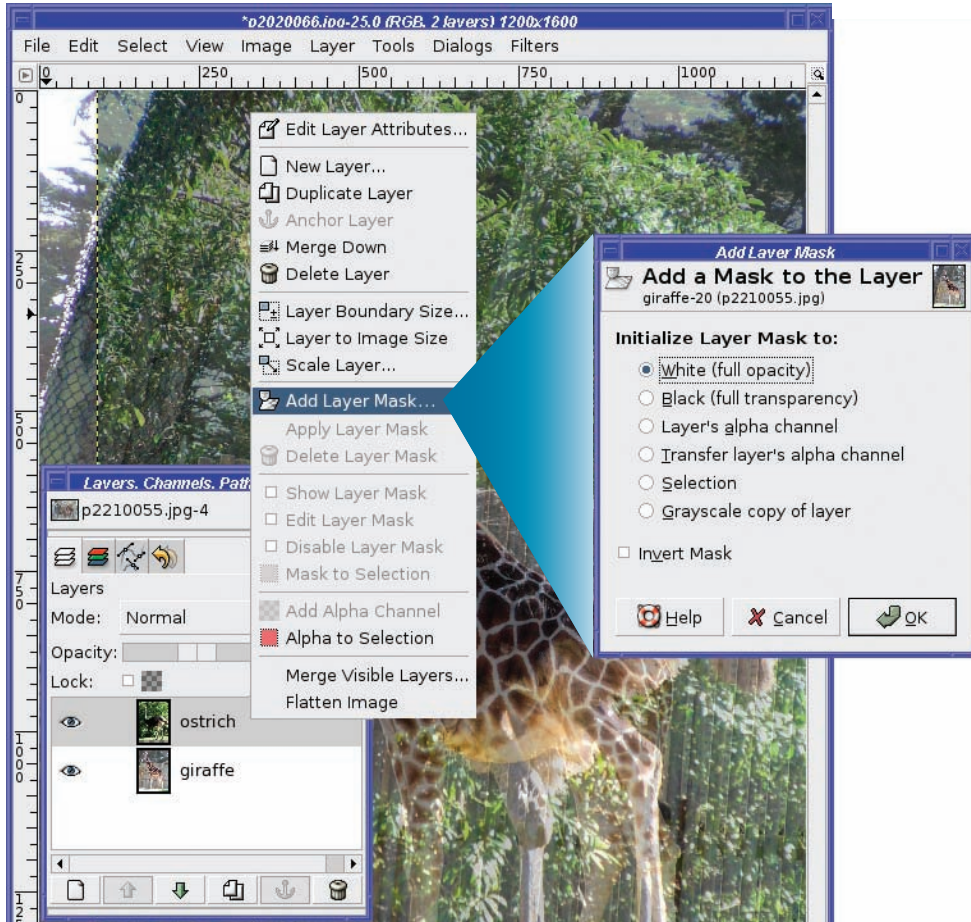
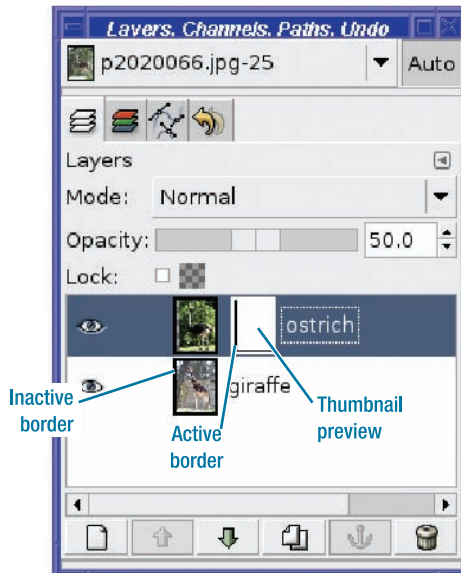


Figure 5-40. Adding a layer mask

The Add Layer Mask dialog lets you specify the color of the new layer mask. The default, *White*, means the full layer remains visible at first (until you change that by painting some black on the mask). This is almost always what you want, but the other options are there if you need them.

The new layer mask shows up as a thumbnail preview in the Layers dialog (Figure 5-41). There are two important points to note about the layer mask preview.



**Figure 5-41.** A preview appears for the new layer mask.

First, it shows you what the layer mask looks like. The new mask is all white, so it looks like a white rectangle. Once you start drawing in the mask, you'll see that the preview will give you an approximate idea of the mask's appearance.

Second, and more important: the border of the preview is white, indicating that it is *active*. That means if you draw in the image, you'll be drawing on the layer mask (just as you drew on the QuickMask or on a selection channel), not on the image itself.

This can be quite confusing. An image normally shows no visible indication of whether there's a layer mask active. When working with layer masks, it's important to keep the Layers dialog visible. Use the preview border colors to tell you whether you're working with the layer or with its mask. The one that has a white border is the one that's active.

You can select the layer at any time by clicking on the layer preview in the Layers dialog. Click again on the mask to select it. Or use *Edit Layer Mask* in the context menu to toggle between the two.

---


**Caution** When a layer mask is selected, drawing isn't the only GIMP operation that will be affected. The Move tool will move the mask, not the layer (if you had the mask lined up with the layer, it won't be lined up any more). Fortunately, Undo fixes minor mistakes like that. Also, saving to some file types will issue warnings if you try to save when you have a layer mask selected. You can cancel the save, go back and select the layer, and then save again.

---



Now try drawing on the layer mask, to see for yourself how it works. Gradients are particularly useful with layer masks. I'm going to keep the head, neck, and back of the ostrich, while making the legs invisible. The ostrich's upper body should blend in to the giraffe's lower body. That means I need the top of the ostrich's layer mask to be white, the bottom black, and a fade (a gradient) in between, right where I want the two images to merge.

With the layer mask still selected, I choose the Gradient tool in the Toolbox and drag a vertical line where I want the fade (Figure 5-42).



**Figure 5-42.** Drag where you want the gradient to be.




It's fairly easy to forget in which direction to draw a gradient, and end up blackening what you wanted to whiten, and vice versa. (Remember, the layer mask preview will show you a thumbnail of what the mask looks like.) Fortunately, that's easy to fix: just draw another gradient in the opposite direction. You don't even need to undo!

Don't forget that holding the Control key down as you drag a gradient constrains it to a multiple of 15 degrees. That's an easy way to get a gradient that's exactly vertical or exactly horizontal.

Once you've drawn on the layer mask, the image doesn't look that different, because of that translucent top layer. It's finally time to slide the opacity of the top layer back to 100%. You could have done that earlier, but without being able to see the bottom layer it would have been difficult to see where to place the gradient.

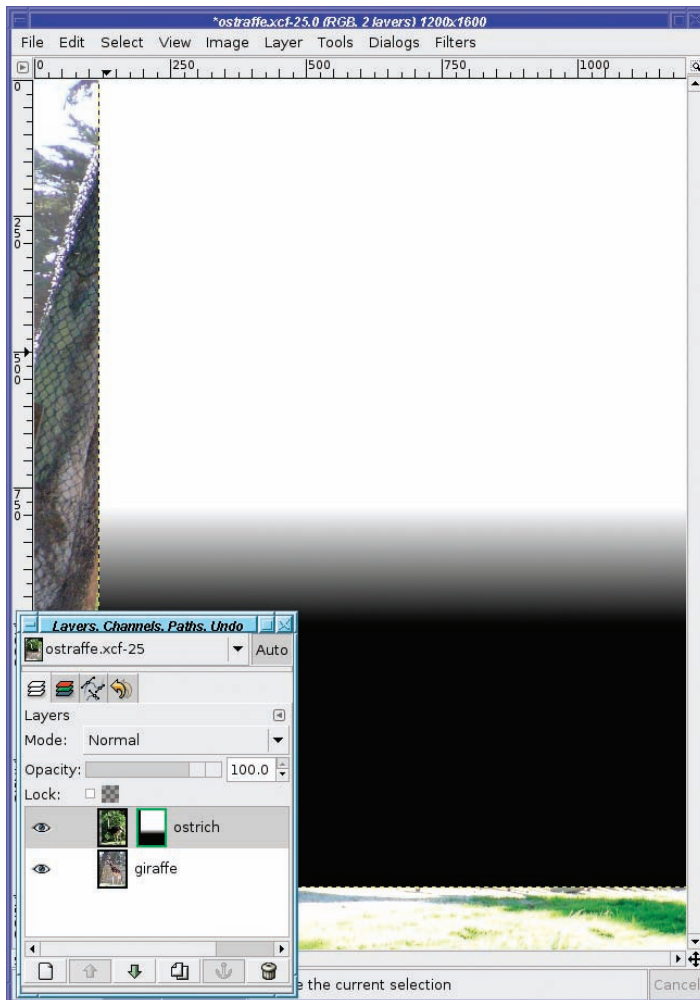
With the top layer made fully opaque again, now the image becomes Figure 5-43.



**Figure 5-43.** A gradient on the layer mask, and top layer back to fully opaque

Notice the layer mask preview has become fuzzy on the bottom: this is the gradient that has been drawn on the layer mask.

To see the layer mask at full size, use *Show Layer Mask* from the right-click menu in the Layers dialog (Figure 5-44). A shortcut for *Show Layer Mask* is to Alt-click (or Alt-Shift-click) on the layer preview; another Alt-click will show the image normally again. Notice that the layer mask's boundary changes to green in the dialog preview. Also notice that the layer mask applies only to the area where the ostrich image is; the edges, from the giraffe photo, are not affected.



**Figure 5-44.** *Show Layer Mask*

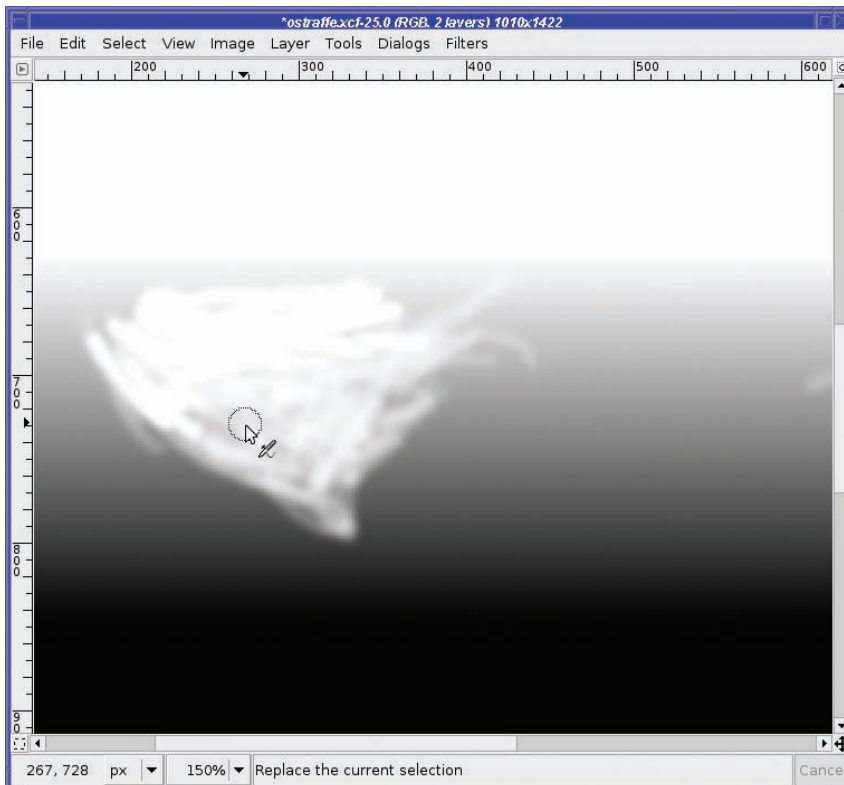
Now play with the image. You may want to change the location of your gradient slightly (by drawing a new gradient higher or lower), or use the Move tool to adjust the relative positions of the two layers. (Don't forget to select the layer, not the mask, before moving.)

When you're satisfied that you have everything aligned just right, it's a good time to crop the image. The idea is to get rid of any overhang from either original that might be noticeable around the borders. For instance, the netting-covered tree on the left side of Figure 5-43 doesn't blend well with the bush to its right.

Now it's time to tune the mask a bit with a little hand-painting. Notice the area under the ostrich's neck where the layer is partially transparent and doesn't blend well into the giraffe's chest? Remember, a mask is just a black-and-white image. You can paint on it with paint tools, just as you painted on the QuickMask and channels earlier in this chapter. Anywhere you paint white will become opaque.


Choose a fuzzy brush with the Paintbrush tool, or use the Airbrush, and paint in the area where you want more of the upper image to show up (Figure 5-45). You may find it helpful to flip back and forth between "Show Layer Mask" mode and normal mode to see what you're painting. A quick way of doing that is by Alt-clicking on the layer mask preview (Alt-Shift-click will also work on systems that use Alt-click for other purposes). The border of the preview will go green to show that you're seeing the layer mask itself, not its effect on the image.

If you ever need to disable a layer mask temporarily, a quick way to do that is to Control-click on the layer preview. The border will turn red, to show the mask is disabled. Of course, you can always show, hide, enable, or disable layer masks through the context menu you access by right-clicking in the Layers dialog.



**Figure 5-45.** Use Show Layer Mask while painting. Use painting tools to touch up areas where the gradient isn't quite enough.

After touching up the layer mask, Figure 5-46 shows the final result.

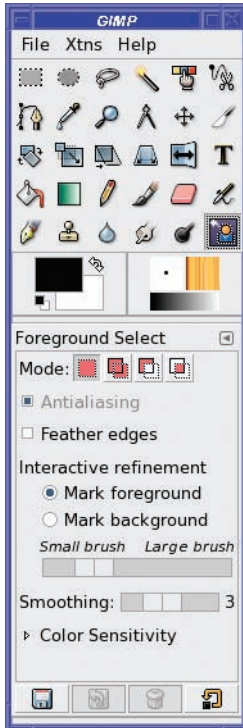


**Figure 5-46.** *The only ostraffe in captivity!*

Of course, the backgrounds from your two images may not match up. You can paint on a layer mask—much as you would on a QuickMask—to hide one background completely and use only the other. Or you can hide the backgrounds on *both* images with layer masks, and then add *another* background layer to place your creation anywhere you want. The choice is yours!

## Extracting Foreground Objects with SIOX

A brand-new tool in GIMP 2.4 is the Foreground Object Extractor (Figure 5-47). It's also called SIOX, which stands for *Simple Interactive Object Extraction*.

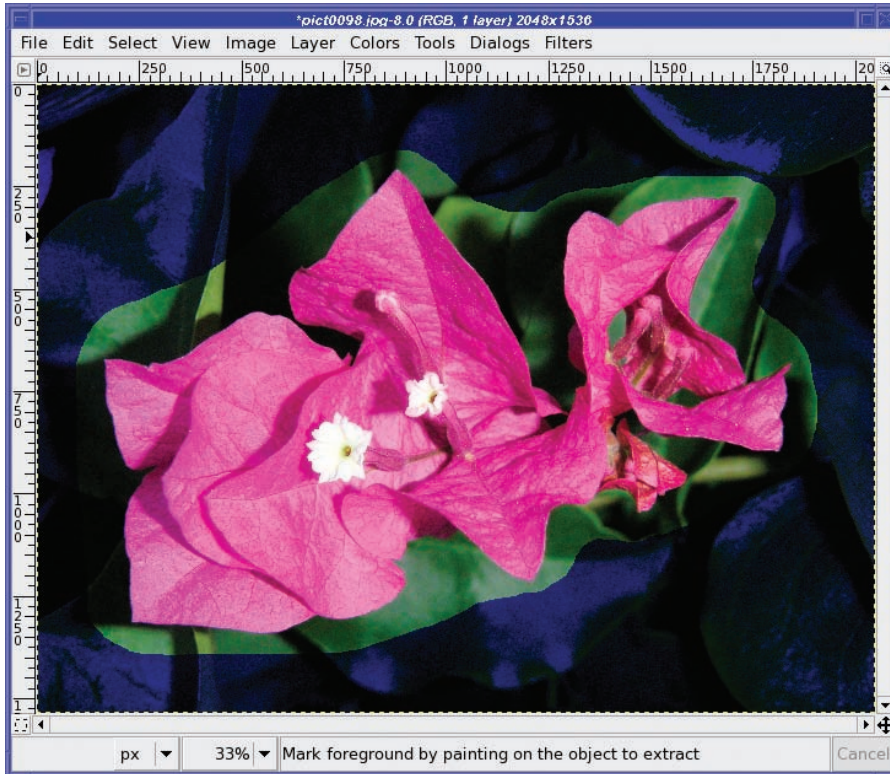


**Figure 5-47.** *The Foreground Object Extractor*

There are two steps to using SIOX. First, draw an outline around and outside the object you want to select. You have to do this in one motion, as if you were using the Free Select tool.




Don't try too hard to get close to the borders of the object. Just make sure that every bit of your object is inside the border. When you're finished, release the mouse button: everything outside the region you selected will be masked off in dark blue (Figure 5-48).



**Figure 5-48.** Define the region of interest.

The second step is marking the foreground. Drag the mouse across the part of the image you *do* want to select, being careful to stay away from parts of the image you *don't* want. The tool will draw a black trail wherever you drag, so you can see what you've already marked (Figure 5-49).






**Figure 5-49.** Drag to define parts of the object you want to select.

You can drag multiple short trails or one long one. I usually find that one long trail works best—sometimes as I add new trails, I lose parts of the image that were previously marked. But each image is a little different, and you'll have to experiment.

When you release the mouse button after drawing a trail, the tool goes to work finding all parts of the region that are similar to the parts you dragged (Figure 5-50). This may take a few seconds, or longer for a large image or a long trail.



**Figure 5-50.** *The tool has done its work.*

Once you're happy with the result, hit the Return key to turn the mask into a selection.

The tool options (back in Figure 5-47) include the usual *Antialiasing* and *Feather* options, and an adjustment for the brush size (you will want to adjust this if you're trying to select very small objects, or to mark thin structures like branches). There's also a *Smoothing* control, which lets you adjust the coarseness of the areas selected, and a *Color Sensitivity* expander that lets you play with how the tool determines what's close enough to the trail you dragged.

The SIOX tool is very young, but it shows a lot of promise. Since it's new, the GIMP developers are still adjusting it, and it's likely to evolve over the next few versions. It doesn't work for everything. I find it works best when you have significant color differences (like a magenta flower against a green background), and not so well for objects that include colors similar to the background (like a brown squirrel perched in a tangle of brown branches).

Try it out. For certain types of images, the tool works very well indeed, and should be a part of your selection arsenal.

## Summary

By now, you should be familiar with a wide range of selection techniques. You're familiar with the marching ants, but you know how to turn them off when they get in the way. You can make selections using color or contiguous regions, Bezier paths, the Intelligent Scissors, or the QuickMask. You can make simple selections into more complex ones by combining several selections in Add or Subtract mode, or by editing them in the Paths tool, the QuickMask, or as a channel mask.

In addition, you've seen how you can change the appearance of an image by using tools on only a selected part of that image, and you've worked with layer masks and seen how you can use gradients in a layer mask to combine two layers smoothly. You can save a selection to a channel or path, and restore earlier selections you saved that way.

Selections are useful in every aspect of image editing, but nowhere are they more useful than when working with digital photographs—especially photos that have something wrong with them that needs to be corrected. In the next chapter, let's take another look at digital photographs, and explore some techniques for touching them up . . . or editing problematic objects out of an otherwise good photo.





# Erasing and Touching Up

In Chapter 2, you learned some basic GIMP operations for improving photographs: the cropping, rescaling, and brightness/contrast tools.

But what if you have a photograph that needs more help than that? I know I sometimes take a photo that comes out just as I'd hoped, except that there's a lamppost growing out of someone's head, or there is an overflowing garbage can right behind the subject. Or maybe that photo of the Grand Canyon would look so much nicer without the guy in Bermuda shorts talking on his cell phone.

In an ideal world, we'd notice these things before taking the picture. We'd be able to move a little and re-frame the photo to eliminate any distractions. But in practice, we all take at least some photos like that. Fortunately, GIMP can help!

This chapter will introduce you to the following concepts:

- Dimming highlights with “Dodge and Burn”
- Smudging blemishes away
- Using the Clone tool for more difficult jobs
- Copying small regions
- Sharpening specific regions using the Convolve tool
- Blurring backgrounds with Gaussian Blur
- Correcting color balance

## Dimming Highlights with Dodge and Burn

I don't know about you, but I take a lot of bad shots of people. Casual “models” like to stand with their backs to the sun, or wear a hat with a visor to keep the sun out of their eyes. I fire away, thinking it'll be a good portrait. But when I look at it later, it's a shot of someone whose face fades into darkness (Figure 6-1).

The answer, as any book on portrait photography will tell you, is to use fill flash: turn on the flash even though it's daylight. It illuminates the face even if the sun is in the wrong place. I've read those books—and I still forget to do it. Do you?



**Figure 6-1.** *Forgot to use fill flash . . . again!*

Fortunately, GIMP has the perfect tool to fix that, along with a host of other exposure errors: the Dodge/Burn tool. Typically, it's used when some small part of the image is too dark or too light.

*Dodge* and *burn* are terms from film photography. When making a print from a negative, some areas may come out too bright, others too dark. Photographic film can record wider *exposure latitude* than can be easily seen by viewing either the negative or the print directly.

The developer turning that negative into a photographic print gets to choose whether to use a long exposure (which will show detail in the light areas, but leave the dark areas too dark) or a shorter exposure (which will make the light areas solid white, but show more detail in the shadows).

That's not always good enough. When professionals make prints, they correct light and dark areas by hand. They place masks over small parts of the image to keep bright areas from becoming too white (*dodging*), and to make sure dark areas get a little lighter (*burning*).

The great photographer Ansel Adams often said that the photographic negative was like a musical score, but the print was the performance. The time he spent making his prints perfect was a big part of what made his photographs so famous.

A digital camera can record at least as much exposure latitude as film can. Your photographs contain a lot more information than you can see just by looking at them. The GIMP's Dodge/Burn tool lets you adjust your digital photographs in the same way that professional photographers adjust their prints. And unlike Ansel Adams' technique, it doesn't require years of experience to learn! Or cost a bunch if you mess up.



## Dodging

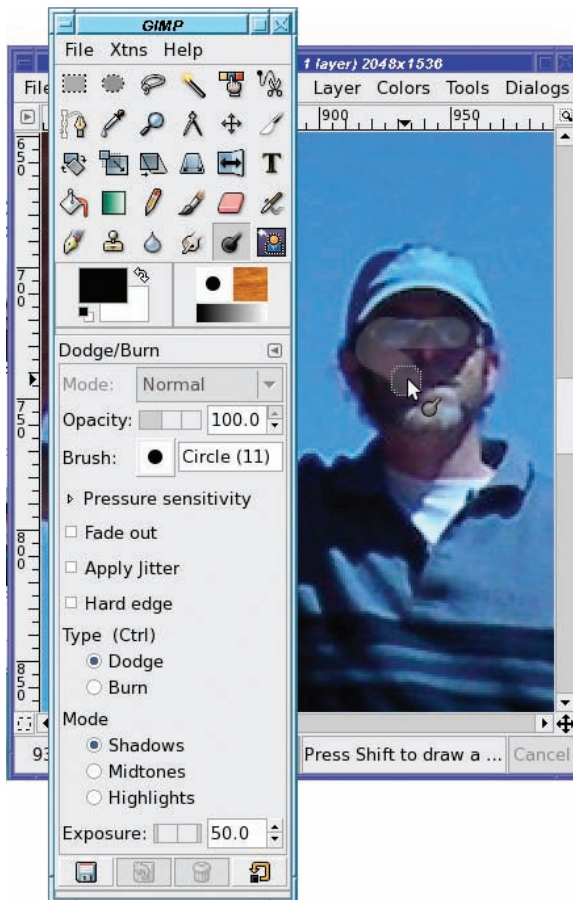
The Dodge/Burn tool starts in Dodge mode by default. It uses a brush—try a fairly large one to begin with, and then change to a smaller or soft-edged brush later.

The Dodge tool lets you choose between three modes: Shadows, Midtones, and Highlights. By default, it's set to Midtones, and that will work for many dodge or burn problems.

However, if you're trying to brighten a very dark shadow area, you may be better off choosing Shadows. Then the Dodge tool will selectively brighten the dark zones much more than other areas. That allows you to be slightly sloppier with your brush technique: if you slide a little outside of the shadow, it won't be very noticeable.

Try that with Midtones and you'll see some of the area that's not in shadow also get brightened. However, Midtones may sometimes bring out a little more detail than Shadows.

Take a look at the Dodge tool's options, as shown in Figure 6-2.



**Figure 6-2.** Use the Dodge tool to lighten areas that are too dark.

*Opacity* makes the tool's effect more subtle—as if the brush were more transparent (the default setting is 100%).

The *Pressure sensitivity* options control how the tool can respond to pressure if you're using a drawing tablet. Mouse users can ignore them.

*Fade out* makes the stroke end after a certain distance. This is one way to ensure that you use short strokes!

*Apply jitter* is a new option in GIMP 2.4 which makes the effect less regular. On some images this may give a more natural effect. The Dodge tool normally makes the edges of the brush slightly fuzzy, even if you're using a hard brush. *Hard edge* turns this off.

Finally, an *Exposure* slider controls the strength of the effect.

Time to get started! Choose Shadows or Midtones depending on how dark the area is, and then start scribbling.

The trick to dodging? Hold the mouse button down and try to cover a contiguous area all in one sweep. The Dodge tool is smart: it won't brighten the same area twice, so if you cross back over an area you've already covered, it won't get too bright. However, if you use separate strokes, you can brighten the same area more each time.

Dodging can turn a failed portrait into a better one. It still won't look as good as if you'd used fill flash or the correct lighting, but it can make the difference between seeing someone's face and not seeing it.

But wait—can't you do all of this by making a feathered selection that covers the shadowed area, then using one of the brightness/contrast tools to make the area brighter?

Sure! In fact, by selecting and using brightness tools you'll have more control: you can use a tool such as Curves or Levels to brighten just the way you want. And if you use the QuickMask to define the selection, you can have full control over *how much* brightening happens and *where*, by making the selection more transparent in areas where you want the effect to be less pronounced.

But all that takes time. Dodging is a shortcut—a quick way to brighten small dark areas, especially areas that are hard to select. It can be a lot quicker—and often it's good enough.


## Burning

Burning is the opposite of dodging—it makes light areas darker. Unfortunately, it's more difficult to use. The problem? In digital photographs, bright areas are often so “blown out” that there really isn't any detail left. The camera's electronics have reached their saturation point, and all they can do is record white.

When you use the Burn tool on areas that are completely white, all it can do is turn them gray, since it has no way to tell what colors or textures should be there.

That means that the Burn tool is best used on fairly small areas, or on areas that are only a little too bright rather than completely white. If you use it on a large white area, you'll end up with a large gray area, which will be surprisingly noticeable. But if you burn small spots of brightness, you can make them detract less from the subject of the photo.

For instance, in Figure 6-3 the glints off the water aren't adding anything to the image. All they do is distract from the subject and lead the eye away from the geese. The Burn tool can help.



**Figure 6-3.** *The glints off the water distract from the geese.*

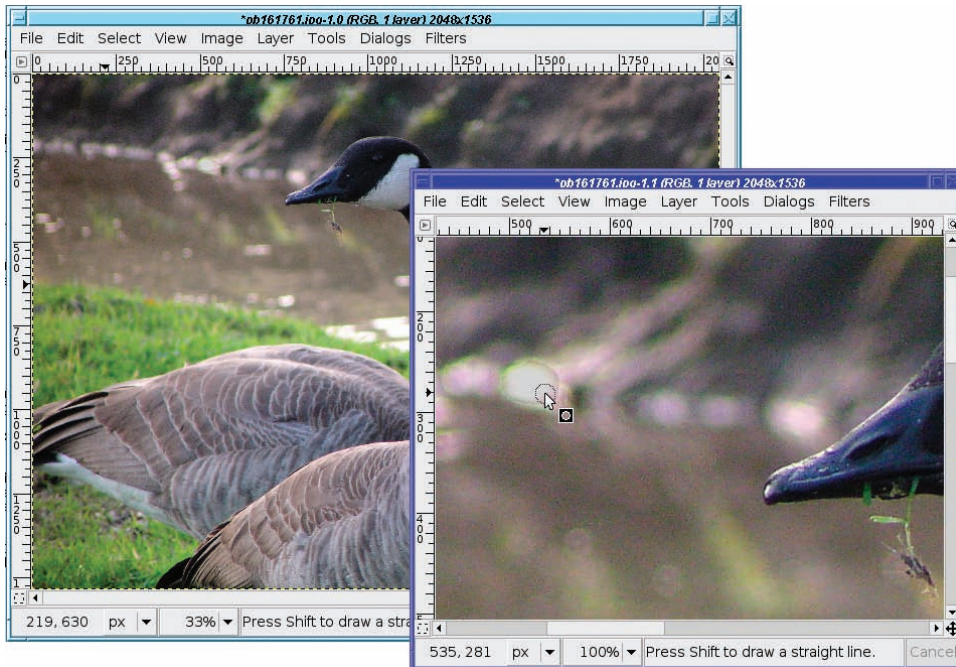
When using Burn, it's especially helpful to open a second view, like you did in Chapter 4 when you were drawing: *View* ► *New View*. That way, you can see the effect of your changes in the full image, even while you work with a nice zoomed-in image.

---

**Note** When burning completely white areas, the mode makes a greater difference than it did when dodging shadow areas. You *must* choose *Highlights*; the other two settings will have no effect on the white areas. (Try it and see!)

---

Now use the Burn tool in the zoomed view, to burn in the bright areas (Figure 6-4). It's hard to watch both views at once, but it turns out there's no need: only the zoomed view will be changed while you're dragging. When you release the mouse button to finish the burn, *then* the other view will change. You can see the full effect of one burn operation. If it isn't what you want, undo and try again.



**Figure 6-4.** Open a new view and zoom in. Use the zoomed view to burn.

Notice that the edges of the burned area in Figure 6-4, burned with the largest hard-edged brush, don't blend in very well with their surroundings. Sometimes this doesn't matter, while other times it does—it depends on the image.

If you burn with a large hard-edged brush and the edges are too harsh in the final image, try using a fuzzy brush and being more careful with your brushstrokes. It takes longer, but you'll get a better end result.

Don't forget that you can burn the same area twice if it's still too bright after one burn. Of course, you can also change the tool's *Exposure* setting if you want a stronger effect everywhere. Be cautious of too much burning, though: it can create flat gray areas that become quite conspicuous in the final image.

### DODGE/BURN AS A DRAWING TOOL


In the GIMP manual, the Dodge/Burn tool is considered a drawing tool because it uses a brush and lets you drag around the image. Although I find it most useful for touching up photographs, Dodge/Burn can come in handy in drawings too.

In particular, you can use it to make shadows along one side of an object you've drawn, for a more three-dimensional look. Remember drawing a dark line down the side of the tree trunk in Chapter 4 to give the trunk depth? You can do that, and better, with the Burn tool.



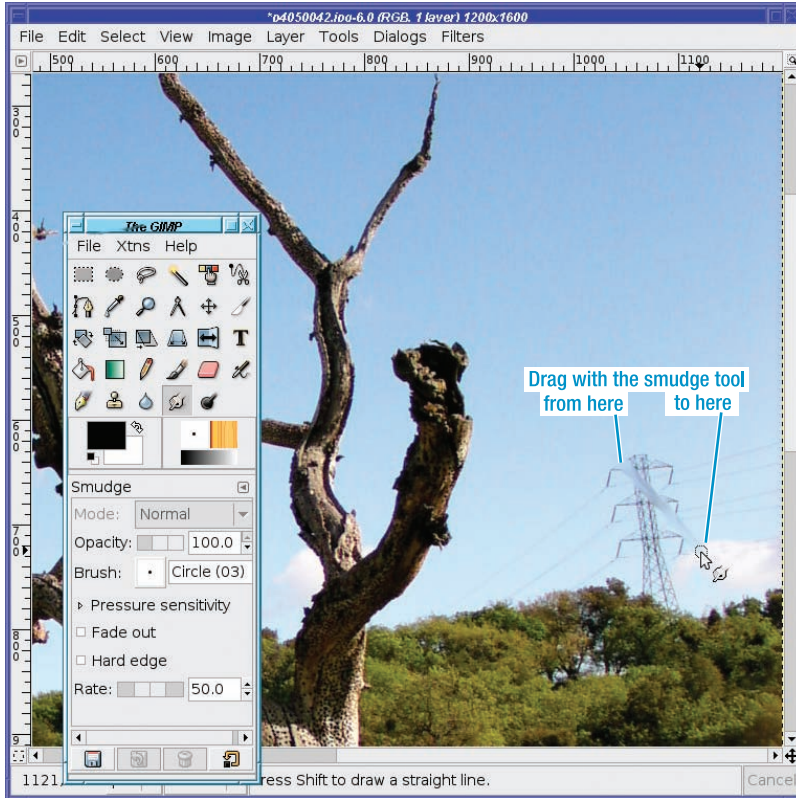
## Smudging Blemishes Away

Many images only need a bit of easy correction. For instance, consider Figure 6-5. The power pylon on the right spoils the tranquil image.



**Figure 6-5.** *Bucolic image spoiled by a power tower*

One way to get rid of simple problems like this is by using the Smudge tool. Its icon in the Toolbox is a finger smudging, and that's exactly what it looks like when you drag the Smudge tool across an image (Figure 6-6). If you ever made finger paintings in kindergarten, the Smudge tool should seem familiar.



**Figure 6-6.** *The Smudge tool*

Notice that the Smudge tool picks up color underneath it, and distributes those colors as you continue to drag.

For this reason, most of the time you'll want to use short strokes when you smudge. Work from the color you want to keep (in this case, the sky) and gradually replace the colors you don't want (the power tower). Drag the sky color in on top of the power tower, and when it starts to turn dark, stop dragging, go back, and start again in the sky. Usually a large brush works best on photographs, because it takes too long to smudge with a small brush. Sometimes you can start with a large hard-edged brush, and then clean up any problems with a smaller fuzzy brush.

Eventually, after many short strokes, the unwanted object is gone (Figure 6-7).






**Figure 6-7.** *The power tower is gone.*

After smudging an object out of a photograph, you may have residual color that you can't seem to get right. If you look closely at Figure 6-7, you may notice a faint darker area in the sky where the power tower used to be. It's very difficult to end up with completely smooth color using the Smudge tool. Often, it's "good enough," and the tool is very simple and fun to use. But if you want cleaner results, fear not! You'll learn several more accurate ways to paint out images.


Smudge is useful in drawing as well as image touch-up. For instance, if you draw a white blob (by filling a rectangle or oval with white) and then smudge in small arcs, you can make new clouds (Figure 6-8). Smudging outward from the white area expands the size of the cloud; smudging inward from the sky into the growing cloud, adds some darker “shadow” highlights that can make a cloud look more realistic.



**Figure 6-8.** Smudging to make larger puffy clouds

The Smudge tool has the same *Opacity*, *Pressure sensitivity*, *Fade out*, and *Hard edge* options as the Dodge tool. In addition, it has a special option of its own: *Rate*.


*Rate* controls the strength of the smudge effect. The default setting of 50% is good for most touch-up work. But you can get some interesting effects by increasing the rate. 100% isn't very useful for image touch-ups, but it gives a very strong smudge with a hard edge and no transparency that can be useful in drawing (Figure 6-9).



**Figure 6-9.** Strange effects from a smudge rate of 100% and a large hard-edged brush. The blue on top of the “tube” is just the sky from above the branch when starting the smudge.

## The Clone Tool, For More Difficult Jobs


The Smudge tool is fun and easy, but if you’ve used it, you’ve probably already hit its limitations. It’s hard to keep control of colors when you’re smudging. More important, smudging smooths out any textures that should be in the original. It’s okay for clear blue skies and other very evenly colored areas, but it wouldn’t work to remove an object on a more complicated background, like the cars in Figure 6-10.



**Figure 6-10.** *I waited and waited for a clear moment, but the cars just kept coming!*

If you try to smudge objects out of areas like this, you'll just make a mess, like in Figure 6-11 (and give yourself a repetitive stress injury from pressing the mouse button too much). This inaccuracy is why finger painting never caught on in professional art circles, even though it was lots of fun to do in kindergarten.





**Figure 6-11.** *Smudging just makes a mess.*

The answer is the Clone tool. The Clone tool lets you duplicate a small area over and over, to paint an object out of a picture.

The first time you try to use it, you'll probably be frustrated. You click and drag, expecting patterns from nearby areas to be copied—just like with Smudge, only smarter. But nothing happens! The cursor shows a crossed-circle (Figure 6-12) saying “No”! Why won't it let you paint?



**Figure 6-12.** *At first, GIMP won't let you clone!*

## Setting the Clone Source

The key to the Clone tool is first selecting the region you want to clone by holding down the Control key while clicking. Once you've selected the region, you can duplicate it over and over by dragging, painting just as if you were using the Paintbrush tool.

This is much more powerful than simply dragging whatever is nearby, like Smudge. You can even Control-click in one image, and then drag in another, cloning parts of the first image into the second.

But there's one other trick to using the Clone tool: it doesn't restrict itself to the pattern that was inside the brush outline when you Control-clicked. If you keep dragging, it will continue copying from the source image. This means the direction and length of your drag is very important.

For instance, suppose that you Control-click to choose a source region in the bushes just to the right of the lead car. Then move down a bit and drag along a line from right to left (Figure 6-13).

Getting the bush at the beginning of the dragged line was great, but you can see the drag got carried away. It went too far left in the original source region, grabbing the car, the center line, and a lot of road.





**Figure 6-13.** *Selecting too much with the Clone tool*

For that reason, choose a large source region if you can (though it's not always possible), and use fairly short strokes. Be careful not to drag too far.

Also, if you're trying to paint a specific object out of the picture, use direction to your advantage. For instance, if you clone from below the car, painting with downward strokes (starting at the top of the car and dragging down) will pick up more road, while if you drag upward you'll just copy the car onto itself.

In the case of the lead car, this technique doesn't work so well: the road ahead of the car is curving, while the road behind it is straight. The best bet is to take a source region from the road behind the car, about halfway between the car and foreground rock, and use very short strokes, being careful about the placement of the road's center line.


As with most touch-up techniques, a fuzzy brush usually works best, unless you're cloning near a sharp boundary like the edge of a rock or tree trunk against the sky.

---

**Tip** When covering an area, particularly a long thin area like this one, using Shift-click is helpful. Remember drawing lines in Chapter 4? Click on a point, and then hold down the Shift key while clicking to draw a line to your next point.

---

Change the source region whenever you need to. You'll probably need to grab from several different source areas (in this case, road, centerline, gray dirt, red dirt, and bushes, as in Figure 6-14).




**Figure 6-14.** Cover the unwanted object(s), changing the source region as many times as you need to.

Finally, it's as cloned as it's going to get. But, as usual, it's not perfect.

## Fine-tuning a Clone Job

You may have noticed in Figure 6-14 that the centerline doesn't go around the bend. That's because the Clone tool copies regions at least as large as the brush. With a large brush, it will always copy a centerline that's pointing in the wrong direction. How can you fix that?


One way is to switch to a much smaller brush, Control-click on the centerline to use it as the source region, and then click repeatedly around the corner, making dots of centerline color (Figure 6-15).



**Figure 6-15.** Using one dot at a time to make the curve in the centerline

Use enough dots to fill in the line completely. Or, if you start feeling too dotted, use the Smudge tool (still with a small brush) to close the gaps between the dots.

Voilà! (Figure 6-16.) I don't know why I kept waiting for a gap between cars to take that photo. Getting rid of cars is easy!



**Figure 6-16.** *Cars, gone!*

## Clone Tool Options


I haven't said much about the Clone tool's options (which were shown in Figure 6-12). That's because most of the time, you won't need to change anything.

The collection of settings should be familiar from other tools you've already used: *Mode*, *Opacity*, *Pressure Sensitivity*, *Fade out*, and *Hard edge*.

In addition, you can clone from one of GIMP's built-in patterns instead of from an image by changing the *Source*.

The only complicated option is *Alignment* (Figure 6-17). The default is *Non-aligned*: for each stroke you make, the pattern GIMP will paint begins at the source origin (the place where you Control-clicked) and proceeds from there.





**Figure 6-17.** Alignment modes in the Clone tool


In *Aligned* mode, the first stroke you make after setting the source will begin from the source origin, just as in *Non-aligned* mode. But for subsequent strokes, GIMP will make each stroke relative to the last, so you can paint an exact copy of the source region as you keep making strokes.

In *Registered* mode, painting will clone from wherever you are painting, rather than from where you chose the source image. This is primarily useful for painting parts of one layer into another layer at the same position.

## Copying Small Regions

Cloning works great for filling in small areas, or for painting patterns. But there are times when it's easier to copy a pattern all at once, in a single block.

For instance, the dune view in Figure 6-18 would look better without the figure standing at the lower-left corner.



**Figure 6-18.** *Dunes, with a person in the way*


You could use the Clone tool to get rid of the person, but there's a faster way: copy a region shaped like the object you want to remove, and then paste it over the object.

Start by making a selection around the object. The selection doesn't have to be perfect—just make sure it's big enough to cover the unwanted object completely. With some feathering, the Lasso tool is usually fine for this job.

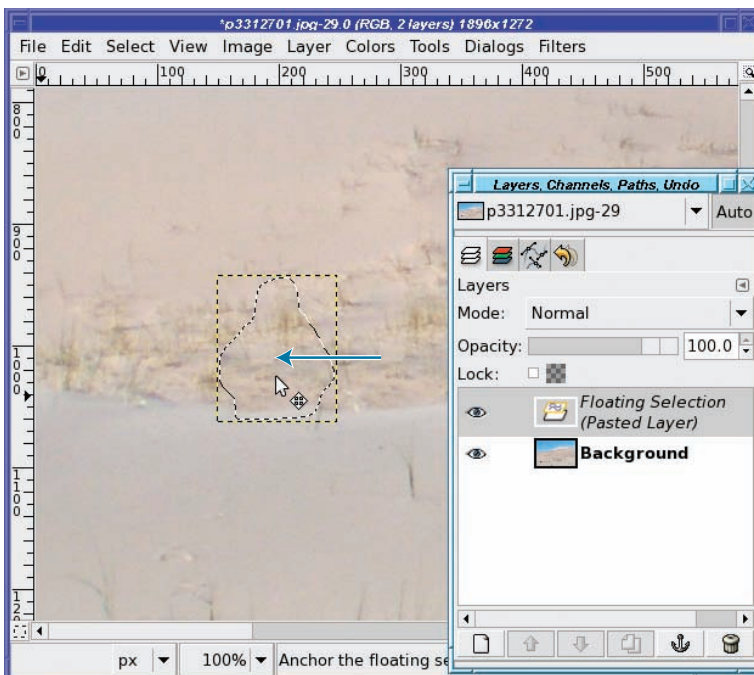
Now move the selection boundary to the area you want to copy. You can use Alt-drag to move the selection boundary (use Shift-Alt-drag if your system uses Alt-drag for something else). Or use the Move tool with *Affect* set to *Transform Selection*, and then just drag from inside the selection.

With the selection boundary in the right place (Figure 6-19), you can copy and then paste. Move the pasted selection back on top of the object you're trying to remove, and then fine-tune its position with the arrow keys until everything meshes (Figure 6-20).





**Figure 6-19.** Move the selection boundary, not its contents.



**Figure 6-20.** Move the pasted selection back over the unwanted object. The selection and layer boundaries are shown.

You will probably want to turn off both the selection boundary (*View* ► *Show Selection*, or *Control+T*) and the layer boundary (*View* ► *Show Layer Boundary*) temporarily, in order to see the edges of the pasted layer better.

That's all there is to it! Quick and easy. If you need to clean up any details around the edges, the Smudge and Clone tools are good for making small corrections.

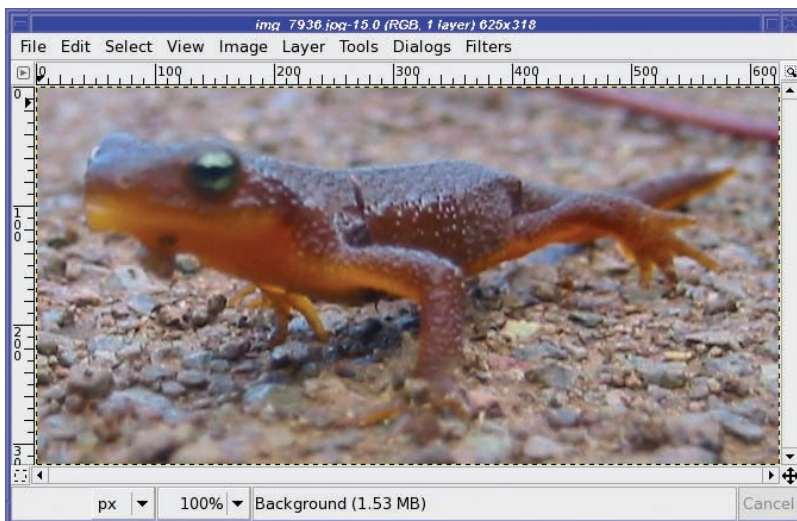
---

**Tip** This technique also works well with “finger-in-front-of-the-lens” mistakes. Not that you'd ever do that!

---

## Sharpening or Blurring Specific Regions: The Convolve Tool

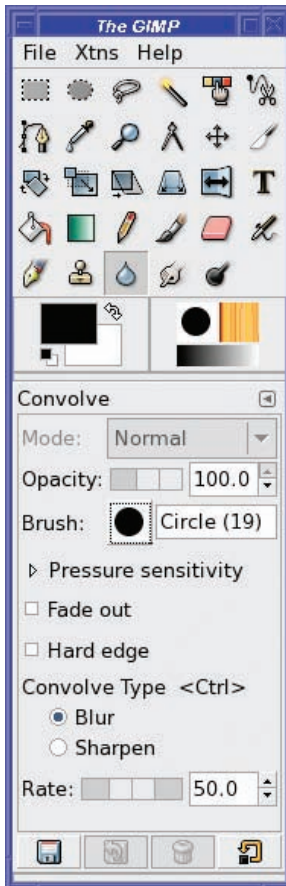
Don't you hate it when you take a picture of something, and it turns out the focus was on the wrong place in the photo (Figure 6-21)?



**Figure 6-21.** *The focus is good on everything except the newt's face!*

You already know about Sharpen and Unsharp Mask from Chapter 2. But if you used those tools on an image like this, the parts that are already sharp would turn grainy. What you really need is something that can sharpen just the newt's face.

That's a job for the Convolve tool (Figure 6-22), more commonly known as Blur or Sharpen.




**Figure 6-22.** The Convolve (Sharpen/Blur) tool

The Convolve tool lets you blur or sharpen small regions of an image just by dragging over them. The more you go over an area, the more it will blur or sharpen.

The tool's options include the usual drawing and touch-up tool options, plus a switch for *Convolve Type* (this is where you choose between *Blur* or *Sharpen*). There's also a *Rate* slider to let you choose the strength of the effect. Since you can drag over the same area repeatedly, *Rate* doesn't matter as much as it does in some tools.

## Sharpening with the Convolve Tool

Dragging the Sharpen tool around the newt's head a little bit produces Figure 6-23.



**Figure 6-23.** *The newt is much sharper now.*

You may notice that the area around the head, and especially the eye, has become a bit grainy. That's what happens when you sharpen something too much (as you may remember from Chapter 2, with Unsharp Mask). If necessary, you can zoom way in and clean up some of that with the Smudge or Clone tools and a very fine brush. By the way, why "Convolve"? Why not just call it "Blur or Sharpen"?

The name comes from *convolution*, the mathematical operation used to blur or sharpen an image.


## Blurring with the Convolve Tool

The flip side of the Convolve tool is Blur. Why would you ever want to make something more blurry?

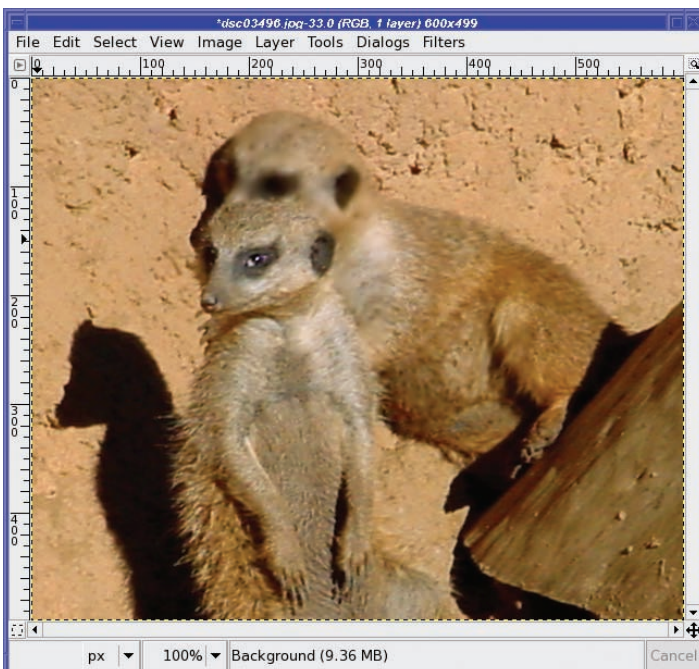
One good reason to blur is to de-emphasize a distracting background, like the second meerkat in Figure 6-24.

A little judicious use of the Blur tool on the rear meerkat's face, and suddenly the front one stands out much more (Figure 6-25). As with Sharpen, you can drag back and forth over the area you want to blur until it's as blurry as you like. (This technique also works well for hiding license plates or other incriminating details.)





**Figure 6-24.** *One meerkat tries to distract your eye from the other.*




**Figure 6-25.** *Now the meerkat in the foreground stands out.*

## DEPTH OF FIELD

The visual property that makes closer objects stand out is related to an effect photographers call *depth of field*.

Some cameras let you control the *aperture* of the lens, or how wide open it is. The *F-stop* is related to the aperture: the bigger the F-stop, the smaller the aperture. At an F-stop of F/2.8, a lens is fairly wide open, whereas at F/11, most of the opening has been closed off, as shown in the figure.



*Large aperture (small F-stop number) means that more light is coming through the lens, but the depth of field will be smaller.*

The curious “division” notation used, F/number, is actually somewhat descriptive. It’s the result of dividing the focal distance of the lens by the current aperture diameter.

A wide open lens lets in a lot of light, so you can take photos in dimly lit rooms or with fast shutter speeds; but the larger the aperture of the lens, the smaller is its depth of field.

Depth of field is the range of distances that are in focus. If your subject is a person five feet away from the camera, the photo might actually have good focus for everything from three feet away to the mountains in the distance (“infinity”)—a large depth of field. Alternately, if the depth of field is very small, perhaps only the range from 4’11” to 5’3” away will be in focus—the person’s nose may be in focus while his ears might not be!

Wouldn’t having everything in focus be a good thing? In fact, it’s often useful to have a shallow depth of field. If your subject is the only object in focus, then the eye is naturally drawn to that subject and away from the rest of the photograph.


Looking back at the meerkats, the original Figure 6-24 shows a large depth of field: both meerkats are in sharp focus. The modified Figure 6-25 simulates the effect of a shallow depth of field: only one meerkat is in focus, and the other one is blurry.

One problem: the photo now looks somewhat strange, because the background meerkat is fuzzy while the wall behind him is sharp. In the section “Blurring Backgrounds with Gaussian Blur,” you’ll see a better way of simulating a shallow depth of field over a wider area.



## Blurring Backgrounds with Gaussian Blur

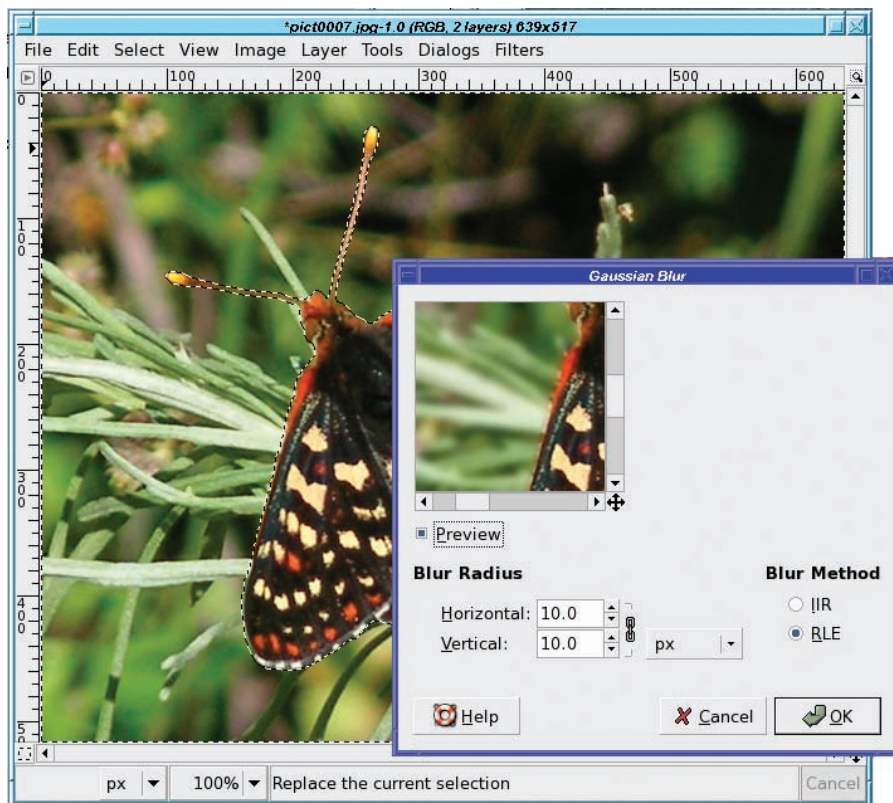
The Blur mode of the Convolve tool was an easy way to make part of a subject stand out. But what if too much of the background is busy and confusing, as in Figure 6-26? You could run the Blur tool over and over the background, but that gets tiresome. There's an easier way.



**Figure 6-26.** *I want to blur the branches on which the butterfly sits.*

First, select the background (everything but the butterfly) using any combination of the techniques from Chapter 5, such as Paths, QuickMask, or Intelligent Scissors.

Then bring up the Blur filter: *Filters* ► *Blur* ► *Gaussian Blur...* (Figure 6-27).



**Figure 6-27.** *The Gaussian Blur filter*

Gaussian Blur can use either of two Blur methods: *IIR* and *RLE*. It really doesn't matter which one you use. They produce the same results, but switching to *RLE* may be slightly faster for images that are computer-generated or have large areas that are all the same color. *IIR*, the default, is faster for photographs.

Adjust the *Blur Radius* (which controls the amount of blur) while watching the preview. Make sure the *Preview* button is checked (it should be, by default). You can scroll around in the preview window if you need to.

As with most effects, the right amount of blur depends on how large (in pixels) your image is, as well as how strong you want the effect to be. Sometimes blurring too much can create an unrealistic, dreamy effect. This could be either good or bad, depending on your goal.


Usually you'll want to keep the *Horizontal* and *Vertical* numbers the same (they're linked together by default), but if you ever need them to be different, you can unlink them.

---

**Tip** Sometimes that tiny preview window is annoying and you want a bigger one. With Blur it's particularly obvious, since it's hard to tell the overall effect of blurring the background from viewing only a small piece of an image. GIMP *can* increase preview size, but it takes some fiddling. See the section on "gtkrc" in Chapter 12 to learn how to change preview size.

---

When you're happy with the blur amount, click OK (Figure 6-28). You'll probably want to toggle the selection off temporarily (*View* ► *Show Selection*) or eliminate it completely (*Select* ► *None*) to see the effect. Make sure you don't have any halos around the selected object!




**Figure 6-28.** *The butterfly stands out more when the grass is blurred.*

## Correcting Color Balance

Cameras don't always record the same colors your eye sees. Photographs taken indoors under incandescent lights often come out looking yellow (though some cameras have settings to compensate for this effect). Fluorescent lights can turn photos blue. Long exposures can have strange effects as well, especially with film cameras.

For instance, Figure 6-29 was a relatively long exposure on color slide film. Although it was a beautiful storm (I was trying to capture lightning, but didn't manage it), the clouds were just ordinary gray, not the purple shade that showed up in the slide.



**Figure 6-29.** Long exposures can create strange color casts, such as this purple sky.

To be honest, I rather like the purple sky in this photograph. But what if I wanted to make the colors more normal?

## Hue-Saturation

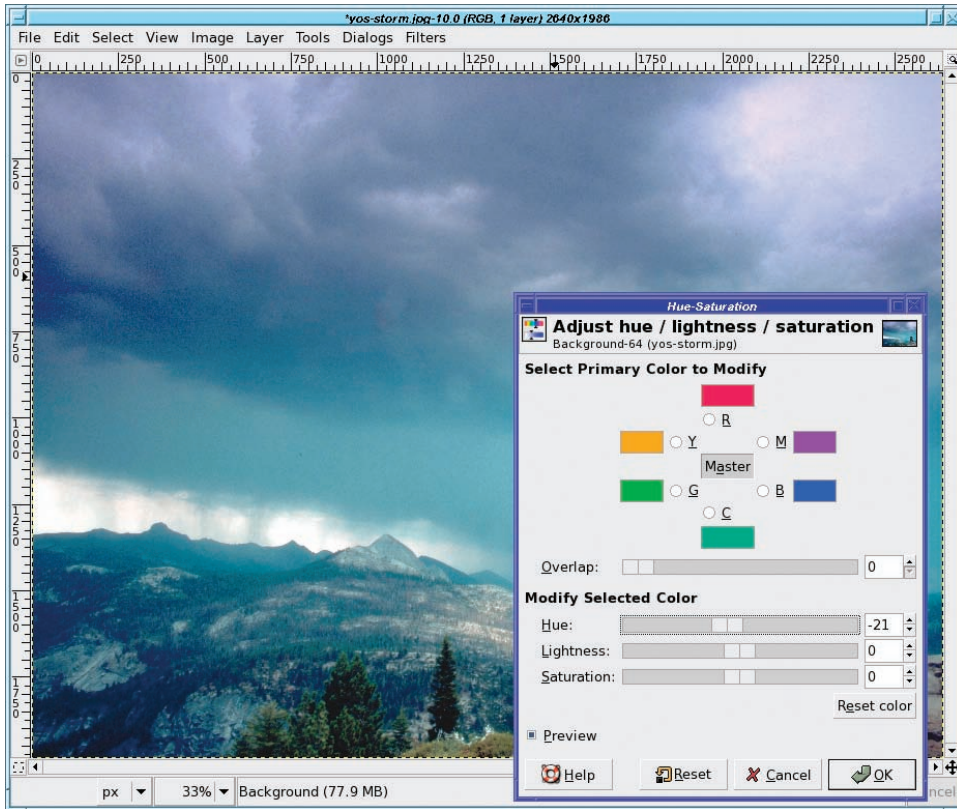
The GIMP has several ways of adjusting color balance. The simplest is **Colors** ► **Hue-Saturation...** (in versions prior to 2.4, find the **Colors** menu inside **Layer**).

You'll learn about the various color models in detail in Chapter 8. For now, all you need to know is that **Hue** represents the color of a pixel: whether it's red, green, blue, or some other shade. **Saturation** is how strong the color is: whether it's intense and bright, or "washed out" and pale. **Value** (shown as **Lightness** in the dialog) is how bright or dark the color is. For correcting color casts, **Hue** is the only slider you need.




You can get an idea of which Hue values represent which color by bringing up the color selector (click on the foreground or background color swatch in GIMP's Toolbox window). Slide the S and V sliders all the way to the right, and look at the colors in the H (Hue) slider.

The effect of adjusting the Hue slider slightly to the left shows up well in Figure 6-30. A subtle change can help eliminate a color cast in a photo.



**Figure 6-30.** A small Hue shift fixes the magenta sky.

However, you usually have to be very sparing with the Hue slider. Move it more than a trifling distance, and you'll have huge color changes in your image, as Figure 6-31 shows.



**Figure 6-31.** Sliding the Hue slightly to the left eliminates a lot of the magenta sky.

## Color Balance

*Layers* > *Colors* > *Color Balance* is another tool for correcting color problems in photographs.

Choose to modify *Shadows*, *Midtones*, or *Highlights*, then slide the sliders in the direction you want to go: in this case, away from Magenta (toward Green) and away from Blue (toward Yellow). Figure 6-32 shows the result.

---

**Caution** A serious limitation of the Color Balance tool is the need to choose *Shadows*, *Midtones*, or *Highlights*. You must choose one, not a combination, and it's not always obvious which is best. In this case, *Highlights* seemed like the obvious choice, because the strongest magenta cast was in the brightest part of the clouds. Wrong! Choosing *Highlights* created a grainy texture in that region, whereas choosing *Midtones* gave nearly the same color balance but without the unwanted texture. When in doubt, use *Midtones*.

---






Figure 6-32. Using the Color Balance tool

## Using Curves or Levels for Balancing Colors

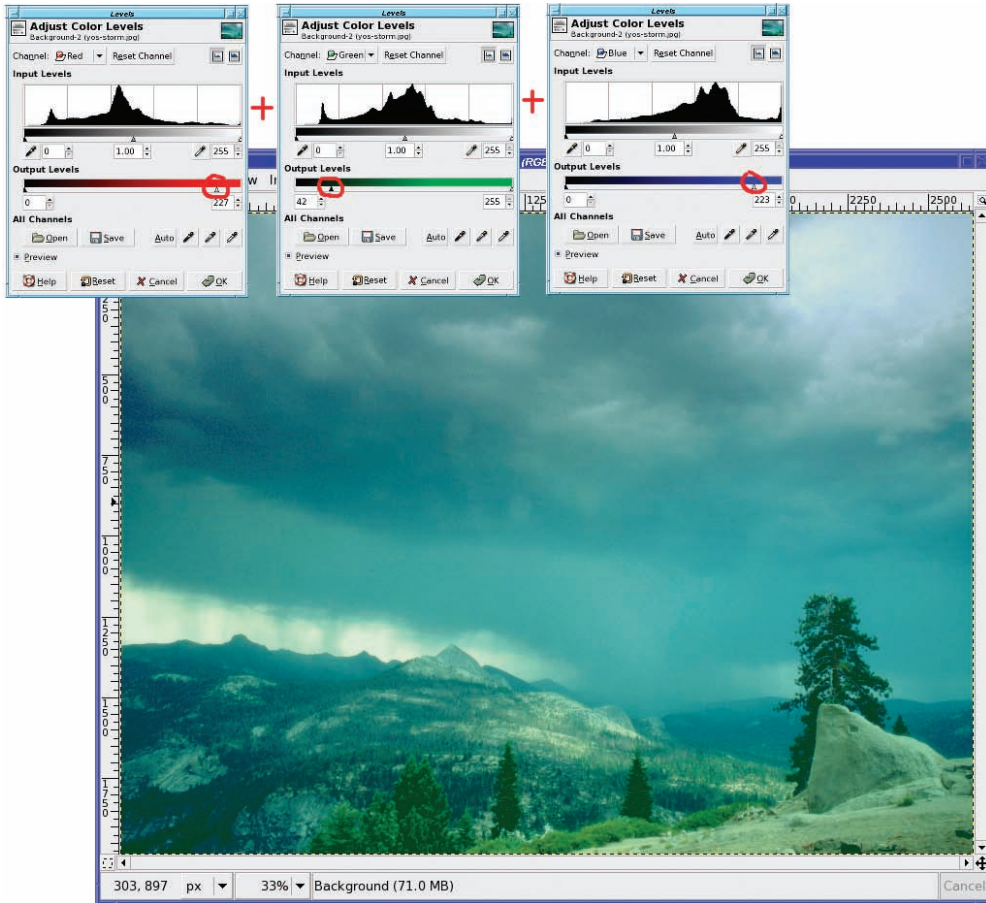
You've already used the Levels and Curves tools for adjusting brightness and contrast. But they're also powerful tools for adjusting color.

The key is the *Channel* menu at the top of the Levels and Curves tools. By default, it's set to *Value*, meaning that the tool will adjust the brightness of the image. But you can also adjust levels or curves of the individual color channels: *Red*, *Green*, or *Blue*.

You can even adjust the transparency, or *Alpha*, channel. However, that is a more advanced technique and isn't generally useful for photographs.

### Levels

Adjust each of the colors independently. Using the Levels tool, I decreased the red and blue highlights, and increased the green shadows (Figure 6-33).

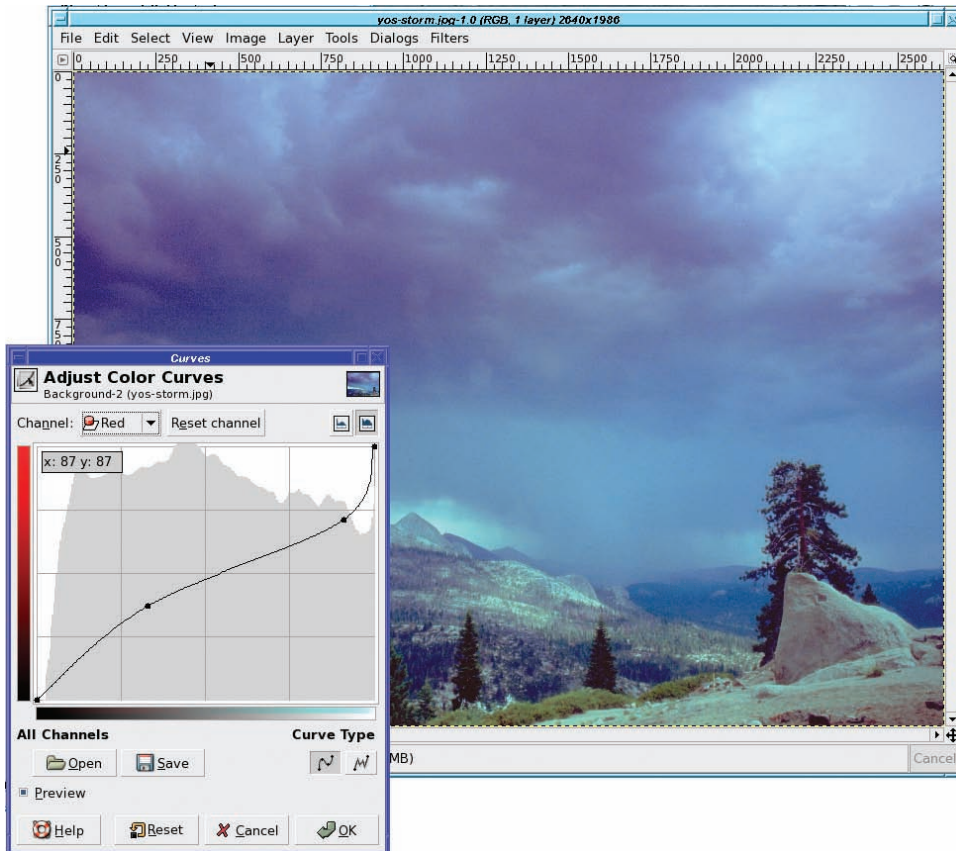


**Figure 6-33.** Using Levels to adjust color balance

## Curves

Curves is even better. It lets you control the amount of each color at each point in the curve.

Figure 6-34 shows the effect of a curve that reduces the amount of red in the brightest areas of the picture, while keeping red the same or even increasing it slightly in darker areas. This gets rid of the red in the clouds without making the rest of the photo blue, a trick that was difficult with Levels. I didn't need to change the other two color channels; making a curve on the red channel was enough.



**Figure 6-34.** *The right curve can make color adjustment easier, even if you adjust only one color channel.*

## Summary

By now you should have a fairly good collection of tools for touching up any problems you might see in your photos.

You can dim bright highlights, and enhance areas that are too dark. You can sharpen or blur specific points, or large regions of an image. You know how to paint unwanted objects out of a photograph or copy patterns from another part of the image (or even from a different image). If a photograph has an odd color cast, you've seen several ways to correct it.

These techniques can keep you busy for quite a while, especially if you have a large collection of photographs that could use some minor touch-up work of one sort or another.

But this is only a tiny sample of the toys the GIMP has to offer! In the next chapter, let's explore some of the other filters, plug-ins, and tools buried within the GIMP.





# Filters and Effects

**T**he GIMP has a huge collection of special effects, tools, and techniques you can use to modify your image or create something brand new.

You've probably already experimented with some of GIMP's plug-ins, but there may be some interesting toys that you haven't seen yet.

In this chapter, I'll take you on a tour of some of the GIMP's filters and special effects. I won't try to cover everything—there's always more to explore—but this chapter should give you a taste of the sorts of effects you can create. You'll learn about the following:

- Image window filters versus Toolbox *Xtns*
- Filters for images
- Scripts to make new images

## Image Window Filters vs. Toolbox *Xtns*

Most of the GIMP's fun special effects are located in two places: the image window's *Filters* menu (plus a few in other menus, such as the *Colors* menu), and the Toolbox window's *Xtns* menu. What's the difference?

Simply, filters operate on an existing image, while items in the Toolbox's *Xtns* menu will create a new image.

The Toolbox options generally create buttons, arrows, bullets and other items usable in web pages or applications, fancy text that you can use for logos or posters, and patterns that you can use as the background to an image or in effects in other images. Plus, there are a few miscellaneous tools to create items, such as a shaded sphere.

Filters can change an image in all sorts of ways. There are straightforward operations you've already used such as blurring, sharpening, or adding drop shadows. There are various artistic filters that can make your photo look like a painting or sketch. There are filters to handle edge detection or noise reduction, and distortion filters that can bend your image like a pretzel.

I'll cover filters first, and then move on to the operations in the Toolbox that can create new images.

## Filters for Images

*Filter* is a general term for anything that can operate on an existing image. Filters will always be found in the menus of an image window, but they aren't all in the *Filters* menu.

First, in GIMP versions prior to 2.4, *scripts* (filters written in GIMP's built-in *Script-Fu* language, or certain other languages such as Perl or Python) appeared in their own top-level menus named for each language, such as *Script-Fu*. In addition, certain plug-ins may also create new menus: for instance, GAP, the GIMP Animation Plug-In, creates a top-level *Video* menu.

Finally, there are some filter operations in other image window menus, such as *Colors* (in pre-2.4 versions, these might appear in *Layer* ► *Colors* or *Image* ► *Colors*).

---

**Tip** Many image filters work only on RGB images, because they need a full range of color. Other filters will only work on a layer with an alpha channel. If you want to use a filter and can't figure out why it's grayed out, check your image's mode (in the title bar of the image window) and make sure it's RGB, not indexed or grayscale. Also check the Layers dialog to see whether the current layer has transparency (right-click on the layer and see if *Add Alpha Channel* is still available). In Chapter 11, you'll learn about how to find out a plug-in's requirements using the Plug-In Browser.

---

## Tools vs. Plug-Ins

Most filters are implemented as plug-ins. If you install a plug-in you've downloaded from the web (you'll learn more about that in Chapter 11), it will most often appear in the *Filters* menu, and will look just like the filters that come installed with GIMP.

The *Colors* menu contains a mixture of *plug-ins* and *tools*. You know about tools: they're operations that have buttons in the Toolbox, like the selection and drawing tools. But did you know that some of the other operations you've used, such as Brightness-Contrast, Levels, and Curves, are also tools?

What's the difference between a tool and a plug-in? Tools can respond to clicks in the image window. When you use a drawing tool, you drag your mouse across the window and leave a trail behind you. When you use a selection tool, you click or drag in the image to indicate where you want the selection to go. Plug-ins can't do that: everything must be done from within the plug-in's own dialog. Equally important, tools can give you a full-size preview of their effect in the image window, while plug-ins can only offer a miniature preview in the plug-in's dialog. Finally, tools display their options in the Tool Options dialog (usually docked below the Toolbox), while plug-ins must use a separate dialog. In menus, tools generally have an icon next to them. Plug-ins usually don't.



There are quite a few tools that don't show up in the Toolbox. You can always get to them from the *Tools* menu in the image window, or from the *File* ► *Dialogs* ► *Tools* dialog in the Toolbox, which gives a list of every tool the GIMP knows about. And you can customize your Toolbox, adding or hiding tools using the “eye” icon in this dialog.

If you forget the difference between tools and plug-ins, don't worry about it. You can use all these filters just fine without remembering which is which.

This chapter won't cover the filters in the *Colors* menu; you've already used some of them, and you'll meet the rest in Chapter 8.

## The Filters Menu

The *Filters* menu is where most GIMP plug-ins live. It's a collection of fun functions that can transform your image in all sorts of different ways. You can spend hours exploring and barely scratch the surface.

The first two items in the *Filters* menu let you repeat the last filter operation. *Repeat* will run whichever filter you ran last (the menu entry will change to tell you what that filter was, e.g., *Repeat “Gaussian Blur”*) with the same settings you used last time. *Re-Show* will bring up the dialog for that filter, so that you can run the same filter but use different parameters this time around. (Not every filter shows up in the *Repeat* menus, though; scripts usually don't.)

*Reset all Filters* lets you back up all the filter dialogs to their default values. Normally, if you run a filter and change any of its parameters from the default, GIMP will remember those settings for the duration of the session. Unfortunately, there's no way to reset just a single plug-in back to its default values, though a tool can be reset with the *Reset to default values* button in the Tool Options dialog.

The rest of the *Filters* menu is divided up into functional areas. Remember, if you're using a version prior to 2.4 and a function isn't where you expect it to be, try looking under the *Script-Fu* menu.

---

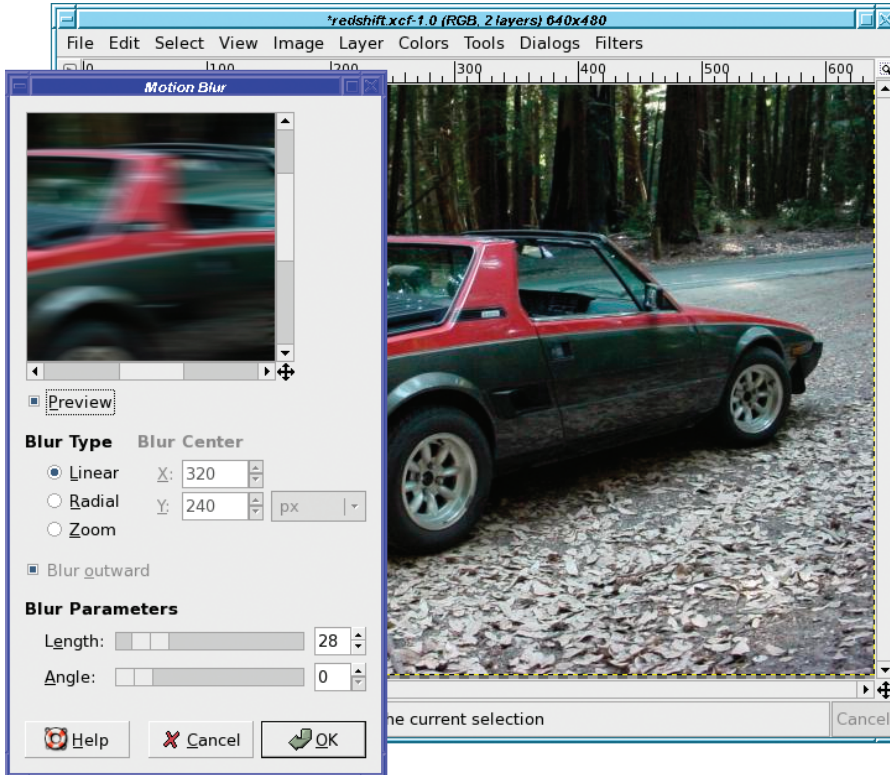
**Tip** In many of these filters, the dialog will show you a preview region. But you may not actually see a preview unless you check the *Preview* checkbox underneath. If you're ever in doubt about whether you're seeing a live preview, look for a box marked *Preview*. With any luck, in some future version of the GIMP it will be possible to say “Always preview.”

---

## Blur

You learned about *Blur* and the very useful *Gaussian Blur* in Chapter 6. But there are several other ways to blur an image.

*Motion Blur* is the most interesting (Figure 7-1). It blurs only in one direction (specified by *Angle*) to give an impression of motion.



**Figure 7-1.** Motion Blur makes a blur in only one direction, giving an impression of motion.

Of course, in addition to giving an impression of motion, this will also make your image look—well, blurry! You can try selecting and blurring just the background, but you have to be fairly precise with the selection. Alternatively, you can get a nice effect by making another layer containing only the foreground object before you blur.

Select the object and eliminate the background (you don't have to be particularly precise with your selection). Paste that on top of the layer that you will blur. After you run the filter, adjust the transparency of the top layer and tweak its position so the sharp layer looks like it's leading, with the blur trailing behind. Experiment and see what sorts of effects you can create.

By default, Motion Blur spreads the image along a line (*Linear*). But it offers two other types of blur. *Radial* is particularly useful for car images such as this one: if you select just a wheel, you can make it look like the wheels are turning. *Zoom* is good for making diagrams showing the expansion of the universe—or just when you want a “zoom” effect for a poster.

For either of these options, you will have to specify the *Blur Center*, the coordinates of the central point for the rotation or zoom. If you know where you want the center, you can easily figure out the coordinates by moving your mouse over the image (without clicking) while watching the status line at the bottom of the image window.

In addition to Motion Blur, the Blur menu offers *Pixelize*, which replaces the image with a mosaic of squares, as though you had blown up a much lower resolution image.

You can also choose *Selective Gaussian Blur*, which is similar to Gaussian Blur except that it tries to keep edges sharp while blurring everything else. Try it and see if you like the effect.

Finally, there's *Tileable Blur*, which gives you a Gaussian-blurred image whose edges can be combined to make a tiled background (you'll learn more about tiling in Chapter 9). It works best if you start with an image that's already tileable.

## Enhance

The *Enhance* filters are used to correct problems in images. The menu contains *Sharpen* and *Unsharp Mask*, which should already be familiar from Chapter 2. But what are those other filters?

*Deinterlace* is intended for images from certain types of video cameras that *interlace*: they record every other line of an image—say, the odd-numbered lines—and then go back and make a second fill in the even lines. If you're capturing a fast-changing scene, this can create still images that look slightly odd. Deinterlace tries to compensate by creating an image using only the even— or only the odd—lines.

*Despeckle* is used to remove small elements of “noise” from an image, particularly specks or small lines caused by dust or scratches on a scanned film image. It can also remove the “moiré” effects sometimes seen in scans of prints. That sounds great; of course in practice, it may also remove detail from your image that you wanted to keep. Use it with caution, but it can be very helpful when you have a noisy scanned image. (A trick that can help is to duplicate the original layer, despeckle it, and then use the Opacity slider in the Layers dialog to mix the despeckled layer with the original.)

You can control the size of the defects Despeckle will look for with the *Radius* parameter. *Black level* and *White level* control how dark or how light a defect has to be for it to be removed. *Adaptive* tries to detect the proper *Radius* setting automatically (with *Adaptive* checked, changing *Radius* by hand may not be effective). *Recursive* makes the filter repeat itself for a stronger effect.

*Destripe* removes vertical lines that are left by some older scanners. Usually, a better solution is to use a decent scanner.

*NL Filter* tries to do everything: it's a combination of smoothing, despeckle, and sharpen in one filter. It always works on the whole image, unlike most filters which will operate only on the selection (if there is one). If you have an image with all these problems, fiddle with the NL Filter settings. Maybe you can make it better.

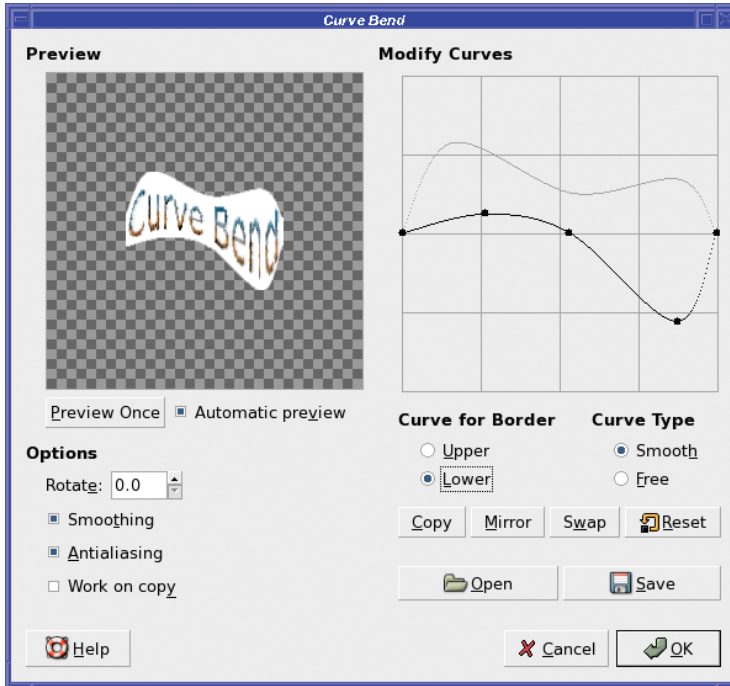
## Distorts

*Distorts* is a fun filter category that offers all sorts of toys. All these filters change the shape of your image in some way.

*Blinds* makes it look like your image is painted on venetian blinds, and you're adjusting the angle of those blinds.

*Curve Bend* is very useful for changing the shape of nearly any layer. It gives you a dialog with a line you can drag to warp the top of your layer to any shape.

But that's not all—you can adjust the bottom of the layer too, by clicking on *Lower* (Figure 7-2). Click *Copy* and you'll get a lower curve that matches the upper curve. *Reset* will make the current curve flat again. *Mirror* will give you a lower curve that's a mirror of the upper. *Swap* will exchange the upper curve with the lower one.



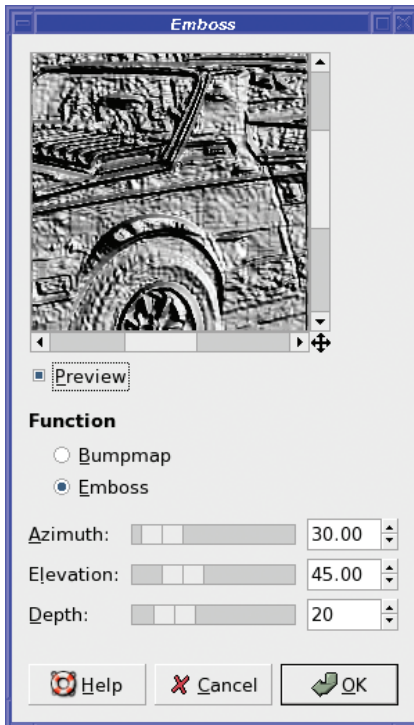
**Figure 7-2.** Select “Lower” to adjust the bottom curve separately.

Choosing *Free* instead of *Smooth* lets you make a spiky curve between your points. *Rotate* lets you turn the result, though it’s not as effective as you might think—you might be better off using GIMP’s free rotation tool (as described in the “Free Rotation” section of Chapter 2) to rotate your curve after you’ve bent it. *Smoothing* and *Antialiasing* do much the same as in drawing tools (most of the time you’ll want to keep them enabled). *Work on copy* creates a new layer in the image and leaves the original layer unchanged. (The new layer will have visibility disabled, so you have to go to the Layers dialog and make it visible.)

Either way, if you started with a layer that was the size of your image, the distorted layer is probably too big to fit. You may have to do an *Image* ► *Fit Canvas to Layers* in order to see all of it. Too bad *Curve Bend* doesn’t do that for you, but don’t let that minor inconvenience stop you from trying a useful distortion tool.

*Emboss* is another nifty filter (Figure 7-3). It makes your image look as though the picture has been embossed onto metal.

The three adjustments are straightforward: *Azimuth* is the direction of the light source in the resulting image, starting from 0 at the right and going counter-clockwise from there (90° is up). *Elevation* is the angle of the light above the image plane (90° means that the light will appear to come from overhead). Play with it and see what it does (usually it makes the effect darker or lighter). *Depth* controls how much relief the emboss will seem to have—that is, how strong the effect is. Finally, selecting *Bumpmap* instead of *Emboss* gives a smoother effect that preserves colors. (It’s actually just regular *Emboss* combined with the original layer in Multiply mode. You’ll see how to create effects like this in Chapter 10.) Don’t confuse this with the extremely important Bump Map plug-in, which you’ll meet later in the Map group of filters.



**Figure 7-3.** The Emboss filter

*Engrave* creates the black-and-white “engraved” look that you sometimes see in old illustrations. It has two options: *Height*, which controls the height of the lines used, and *Limit line width*, which keeps the lines from bleeding together completely in the dark areas (such as in the car’s shadow).

Engrave requires that the image already have an alpha channel. If it’s grayed out, try *Layer* ► *Transparency* ► *Add Alpha Channel*.

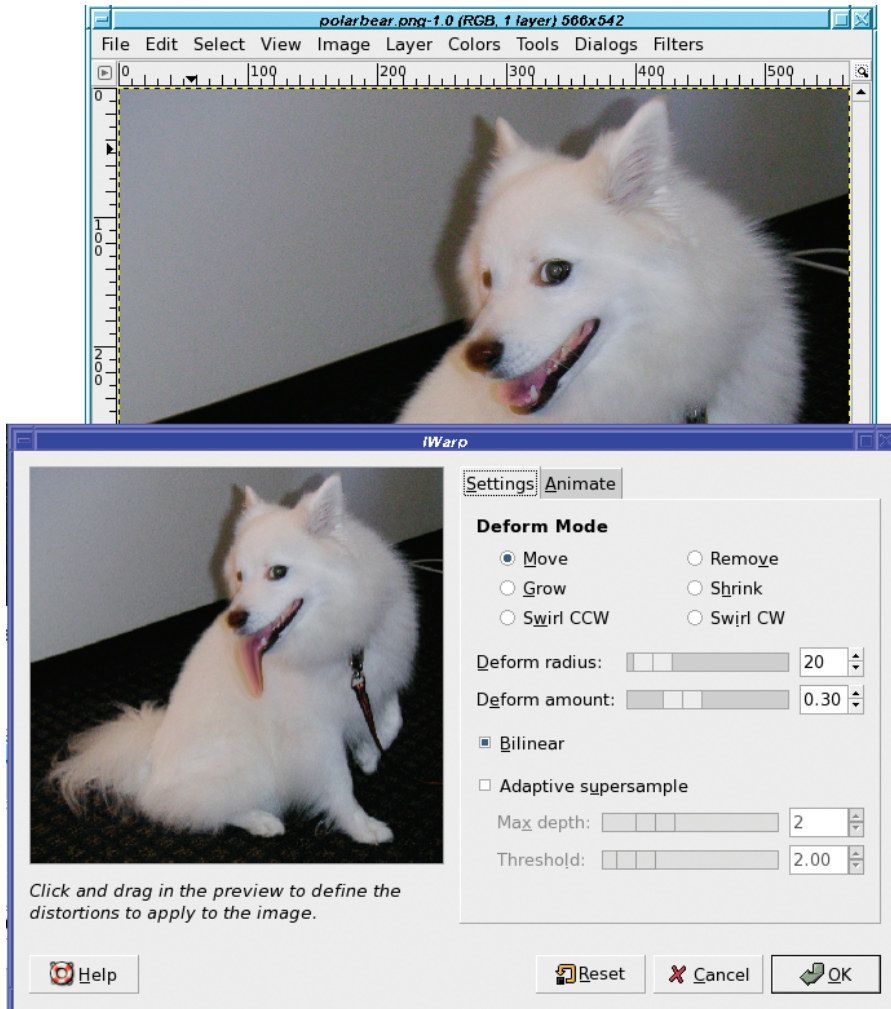
*Erase Every Other Row* (previously located under *Script-Fu* ► *Alchemy*) does just that. It can be useful for certain effects, such as creating the illusion of a television screen, or combining an image with modified versions of itself.

---

**Caution** If the result looks odd to you, check your zoom level. If you’re zoomed less than 100%, you may not see the effect properly. It also may not work right if the layer goes outside the boundaries of the image.

---

*Warp*, the Interactive Warp filter, is one of the GIMP’s best toys. Occasionally, it’s even useful. It brings up a dialog like Figure 7-4.



**Figure 7-4.** Use IWarP to make a dog's tongue longer.

Click and drag in the preview area of the dialog. Parts of the image will drag along with your mouse, “stretching” the image as though it was printed on a rubber sheet and you were deforming the rubber. Use short, repetitive strokes.

You can use IWarP in all sorts of creative ways: you can change animals into other animals, chubby cheeks into gaunt ones, or a human into a Vulcan. It can also be useful as part of other transformations, such as the panoramic images you’ll make in Chapter 10.

IWarP has quite a few *Deform Modes*. The default, *Move*, stretches the image wherever you drag the mouse.

*Remove* lets you remove the distortion in one area of the image. There’s no Undo for the strokes you make inside IWarP, so if you go too far in one place, but you like the changes you’ve made everywhere else, a quick switch to *Remove* mode might help.



*Grow* inflates a round area, as though you placed a magnifying lens only over the area where you're dragging. *Shrink* does the opposite.

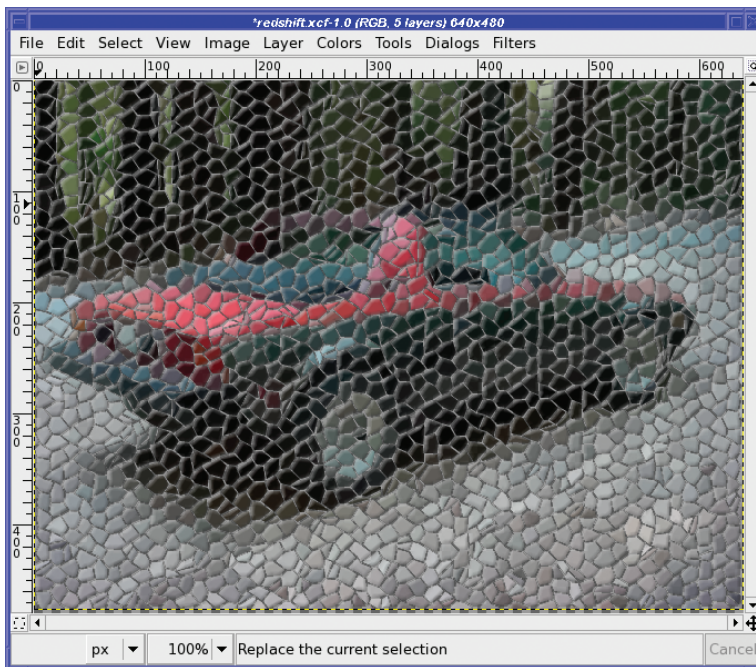
The *Swirl* options mix up the image under the mouse as you drag—as though you were dragging a cake beater through the image. You can make neat paisley swirls out of any picture.

*Deform radius* is like a brush size: it controls how many pixels to either side of the mouse position will be distorted. *Deform amount* controls the strength of the effect.

*Bilinear* makes the effect smoother (it's selected by default), while *Adaptive supersample* and its sub-options *Max depth* and *Threshold* can smooth the effect even more, at the cost of extra CPU work.

There is one problem. You have to make your changes in the small preview in the dialog. Since there's no image-sized preview, it's often difficult to tell when you've gone too far. But don't let that stop you from having fun with one of GIMP's most enjoyable distortion filters.

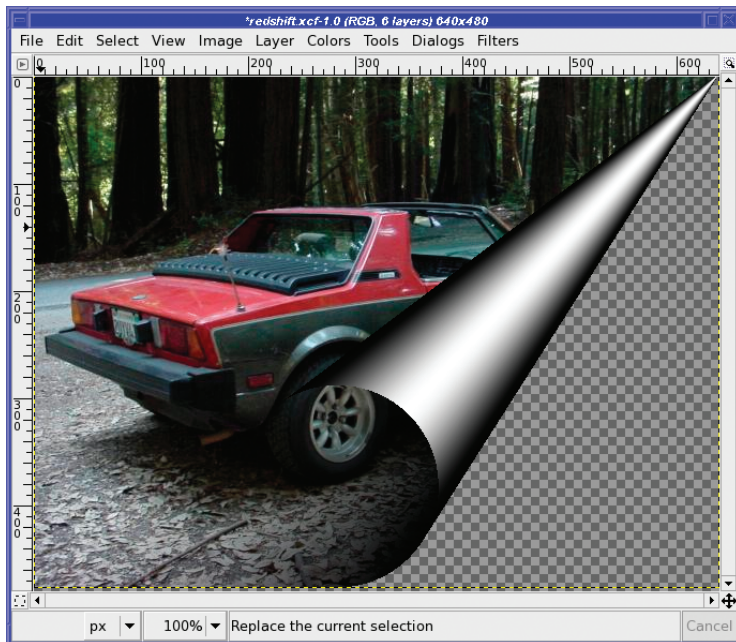
The *Mosaic* filter does exactly what you'd expect: it creates a mosaic out of your image, as if made with random tiles (Figure 7-5). It has lots of options controlling the tile shape size, height, spacing, and surface.



**Figure 7-5.** *The Mosaic filter*

The *Newsprint* filter creates an effect like zooming way in on a newspaper photo, with the individual ink colors visible.

*Pagecurl* gives the very useful effect shown in Figure 7-6. You can control which part of the page is curled and what colors are used for the page. If you want less of the page to curl, try selecting a rectangle in the corner where you want the curl. The size of your selection will control the amount of curl.



**Figure 7-6.** *The Page Curl filter*

The *Polar Coords* filter wraps an image around a point at the top or bottom center. It can also take an image that is wrapped in this way and unwrap it. This is useful for looping text around in a circle (Figure 7-7). Images with straight lines in them can turn into nifty spiral patterns.



**Figure 7-7.** *Polar Coords can wrap text in a circle.*

The *Ripple* filter makes an image look as though it's being viewed through water waves. The options are fairly straightforward and fun to play with, though the *Waves* filter gives a better water effect. The *Border* options control what happens at the edges of the image.

*Shift* moves each row by a random amount, either vertically or horizontally, creating an interesting ripple effect which is quite different from the one *Ripple* makes.

*Value Propagate* is tricky. It finds areas of your image that are nearly black or nearly white (your choice) and expands them. You can set the amount by which areas will be expanded, and how close they have to be to black or white. You can also use it to expand colored areas by selecting *Only foreground* or *Only background*. You can even expand or contract areas of transparency.

Still, even with all these options, you'll probably only want it in very specialized circumstances. It can be useful when working with layer masks.

The *Video* filter applies various patterns intended to make the image look as though it's being viewed on a television screen or various other types of low-resolution device.

Curiously, unlike most filters, it doesn't have a real preview area; instead, it shows a picture of a fox with pre-generated examples of what each pattern should do.


*Waves* is similar to *Ripple* except that the waves are concentric around the center of the image, as if you'd thrown a rock into a swimming pool. It's a nice effect. You can adjust three parameters controlling the size of the waves: *Amplitude*, *Phase*, and *Wavelength*.

*Smear* controls what happens at the edges, much like *Border* did in the *Ripple* plug-in.

*Reflective* makes the waves bounce off the sides of the image and back in, where they interfere with waves coming from the center. It works best if you reduce the amplitude and increase the wavelength. Try it and see.

*Whirl and Pinch* is useful for several types of "funhouse mirror" effects. Its *Whirl* mode grabs the image from the center and stirs it like clothes in a washing machine. You control the amount with the *Whirl angle* slider.

*Pinch* of a negative amount bloats the image as if it's blown up like a balloon from the center of the image (Figure 7-8). A positive pinch does the opposite, compressing the center of the image as though all the pixels were sliding down a drain in the center.



**Figure 7-8.** Left: *whirl* (no pinch). Right: *negative pinch* (no whirl).

**Tip** Pinch can sometimes be used to compensate for lens distortion, though there are plug-ins available on the web that provide more specific anti-lens distortion effects.

*Radius* controls how much of the image will be affected, with “2” meaning the entire image. This is deceptive, though, since the effect is applied smoothly and will usually look like it’s affecting more of the image than it actually is.

The *Wind* filter gives a nice effect similar to Motion Blur, but it’s much more specific. It finds dark or light lines in the image and makes “streaks” trailing in only one direction.


*Strength* controls the amount of wind, while *Threshold* controls how many lines are drawn. *Blast* creates thicker lines than *Wind*, making for a less subtle (and often less effective) result. You can control the *Direction* from which the wind seems to blow. You can also control the *Edge Affected*—*Leading*, *Trailing*, or *Both*—but you may need the opposite setting from the one you expect.

## Light and Shadow

The GIMP offers quite a few filters relating to lighting and shadows. GIMP 2.4 collects all the lighting plug-ins in a single menu, *Filters* ▶ *Light and Shadow*. In earlier GIMP versions, they were scattered in four places: *Filters* ▶ *Light Effects*, *Filters* ▶ *Glass Effects*, *Script-Fu* ▶ *Shadow*, and *Python-Fu* ▶ *Effects*.

## Flares and Sparkles

The *FlareFX* filter is a way to put a sparkle into your picture (Figure 7-9). It has no options to speak of; just click or drag in the preview to set the flare where you want it.



**Figure 7-9.** *FlareFX* adds a bright reflection plus lens flare effects.

*Gradient Flare* is similar, but lets you choose from a selection of different flare effects in the *Selector* tab. The *Settings* tab lets you change all sorts of parameters to adjust the flare exactly the way you want it.

*Lighting Effects* is even more general. It lets you add a beam of light in many different ways.

This plug-in has more options than there is space to cover them, but here are some tips. In the *Light* tab are three lighting *Types*. *None* always begins the light in the upper-left. *Point* displays a blue point somewhere in the preview; you can drag it around, and you should see a beam of light passing through the blue dot to the center of the image. *Directional* is similar to *Point*, but offers a bit more control. On the other hand, with *Directional* it's easy to lose the point as it slips off the edge of the image. The *Bump Map* and *Environment Map* options will make more sense after you've worked with Bump Map (in *Filters* ► *Map*, and covered later in this chapter).

The *Sparkle* filter is a bit similar to the two flare plug-ins, but instead of adding one big flare, it adds small sparkles anywhere that's lighter than a threshold you can set with *Luminosity threshold*. *Flare intensity* and the various *Spike* options control the way each sparkle is drawn. *Opacity* makes the sparkles translucent, making the effect a bit more subtle (higher numbers indicate more transparency, the opposite of what you would expect from the name).

Sparkle's *Random Hue* and *Random Saturation* allow for sparkles of different colors, but the effect is very subtle unless you also increase *Opacity*. *Inverse* puts dark sparkles on the dark areas. *Foreground color* and *Background color* can control the color of the sparkles; as with the *Random* options, the effect is much more pronounced if you add some opacity.

## Shadows


You've already used *Drop Shadow*. *Drop Shadow and Bevel* also makes a drop shadow, but in addition, it *bevels* the layer, sculpting its edges to make it look more three-dimensional. This is primarily useful for text, especially text in light colors (it won't show up in black text since the bevel shadow will be invisible).

The *Perspective* filter makes a shadow tilted at an angle to show perspective. Unfortunately, the plug-in has a bug in most GIMP 2 versions: the shadow appears in a box which prevents you from seeing most of it, so it isn't very useful. With any luck, the problem will be fixed by the time you install GIMP 2.4. If not, you'll learn how to make this sort of shadow yourself in Chapter 9.

## Lens and Glass Effects

*Xach effect* is meant to give the illusion of a translucent object (perhaps made of rough-ground glass) raised above its surroundings. It makes the most sense when you use it on a layer containing a few straightforward objects—for instance, letters in a text layer—above a separate background layer (Figure 7-10). The layer that is active in the Layers dialog when you run the effect is the one that will be used for the shape. Alternately, it will use a selection, if you have one: you can create nice Xach effects by using the Free Select tool to make a selection of any shape you wish to use.





**Figure 7-10.** Xach effect

Named for Xach Beane, a GIMP contributor who came up with the idea, Xach effect creates several new layers, leaving the existing layers untouched. If you use a foreground layer rather than a selection outline for your shape, you'll probably want to turn off visibility on the foreground layer.

*Apply Lens* warps the image in the way that a wide-angle lens might. You can get a similar effect by pinching with a negative number in *Whirl and Pinch*, but *Apply Lens* is more specific and has only one parameter, the index of refraction of the lens being simulated.

*Glass Tile* makes the image look as though it's being viewed through a wall made of many glass bricks, each of which refracts light a little differently.

## Noise Filters

These filters add noise to an image. Why on earth would you want to do that? You'll see as you explore them.

*HSV Noise* (previously called *Scatter HSV*) creates colorful noise. It has four parameters. *Hue* controls the average color of the generated noise. *Saturation* and *Value* control the saturation and brightness of the noise. The other parameter, *Holdness*, controls how much the color of the image is allowed to change. A low value of *Holdness* will preserve more of the original colors than a high value.



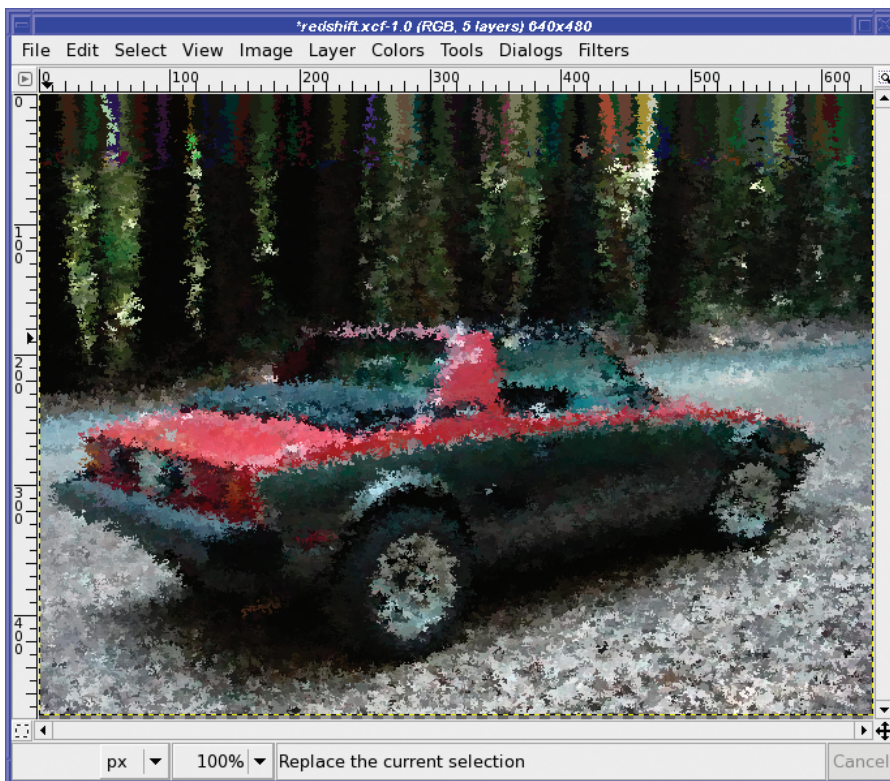
*Hurl* changes a random selection of pixels to random colors. *Randomization* is the percent of pixels to be changed; at 100%, the entire image will be replaced by random noise. *Repeat* runs the effect repeatedly, but you're usually better off just using a higher value for *Randomization*.

*Pick* replaces each affected pixel with a value randomly chosen from nearby pixels. So the image won't change nearly as much as it would with most of the other noise filters. Even with *Randomization* set to 100%, *Pick* can create an image that's quite recognizable and interesting.

*RGB Noise* (known in earlier GIMP versions as *Scatter RGB*) creates very natural-looking colored noise. You can set each color independently, or with *Independent RGB* turned off, you can slide the three sliders together.

RGB Noise can be used on a grayscale image as well; in that case it will show only a single slider.

*Slur* is a bit more fun: it randomly moves pixels downward, as though they were melting. That's not immediately obvious when you run it once, but try setting *Repeat* to a large number like 50 (Figure 7-11).




**Figure 7-11.** *Slur* with a repeat of 65 and randomization of 100%.

*Spread* is somewhat like *Pick*, except that it replaces each affected pixel with one chosen at random from somewhere nearby in the image (how near it must be is controlled by the *Spread Amount* parameters). It makes the image look stippled, as though drawn with dots from blunt-tipped colored marking pens.

## Edge-Detection Filters

The GIMP comes with five edge-detection filters (though there are other ways to detect edges: you'll learn a few in Chapter 10). They're all a bit different.

*Difference of Gaussians* gives a subtle pastel effect by blurring the image with two different radii, and then subtracting one from the other (Figure 7-12). You can tune it by adjusting the two radii.




**Figure 7-12.** The original image, and after using *Difference of Gaussians*

*Edge* is a real workhorse, and includes six different edge-detection algorithms, each a little different. It also lets you tune its sensitivity with the *Amount* slider. Most of its algorithms give a wild neon effect. The more colors in the original image, the better the neon effect will be. *Wrap*, *Smear*, and *Black* are options controlling the effect at the edges of the image, in case you're trying to make a tile that can be used as a repeated background.

*Neon* also gives a neon look, with fewer options and a somewhat different effect (Figure 7-13).

*Laplace* has no options at all. It simply replaces your image with a line drawing, everything else being transparent. Short and sweet! *Sobel* gives a result somewhat like Laplace, but it gives you a chance to tune the algorithm slightly, with a preview. Sobel is a bit more generous in finding lines than Laplace is (Figure 7-14).



**Figure 7-13.** Two ways of making a neon sign: the Edge plug-in with the Sobel algorithm versus the Neon plug-in



**Figure 7-14.** Laplace and Sobel are two ways of finding the edges in an image.

## “Generic” Filters

The *Generic* category is a catch-all for filters that can be used in many different ways. It currently contains only three filters.

*Convolution Matrix* is a general way for people working on mathematical image-processing algorithms to test those algorithms in the GIMP. Under the hood, many of the GIMP’s plug-ins use a convolution matrix. If you don’t know what that is, you may not find this plug-in very interesting, though of course you can try typing random numbers into the matrix to see what happens. You may occasionally see convolution matrix numbers on a web page describing a particular effect, in which case you can copy them. Be sure to enable at least one channel in the list, or the convolution matrix won’t have any effect.

*Dilate* enhances dark areas of the image, reducing the size of light areas. *Erode* does the opposite, enhancing light areas at the expense of dark areas. These filters are especially useful on text when you want to make letters thinner or fatter.

## Combine

The *Combine* group is another catch-all, and contains only two plug-ins.

*Filmstrip* combines the layers in the image into a single image that looks like a strip of 35mm film (Figure 7-15).



**Figure 7-15.** *The Filmstrip plug-in*

It’s most interesting when you use a set of similar but related images—perhaps different photographs of the same subject, or scenes from a vacation. You can load several files as layers in the same image by using *File* ► *Open as Layer*. (To open multiple images at once, use Shift-click or Control-click on the file list in the Open dialog.)

---

**Caution** The *Filmstrip* plug-in insists that all the layers in the image be the same size. If they aren’t, it may crash or otherwise misbehave. Fortunately, when a plug-in crashes, it doesn’t take GIMP down with it, though you might be wise to save any open images after any plug-in crashes.

---

*Depth Merge* combines two images in a complicated way, according to two other images, called *depth maps*. All four images *must* be the same size (the plug-in will only let you choose depth maps the same size as the original images).

*Source 1* and *Source 2* are the two images you want to combine. They can be any layer from any image currently open in the GIMP (as long as they’re the same size, of course).



The two depth maps specify which part of each image will be most prominent in the final image. It's best if the images contain only blacks, grays, and whites; color is meaningless in a depth map.

Depth Merge compares the two depth maps, and whichever one is darker “wins.” Wherever a source image's depth map is black, you will see that source image. Where it's white and the other source image's map is black, you'll see the other source. To try out Depth Merge for the first time, try using two gradients: one that fades from white on the left to black on the right, and a second that is the reverse of the first.

Depth Merge is another way of combining images in a way similar to the ostrich and giraffe project in Chapter 5. You may find it quicker than setting up the images yourself, and it can be useful for combining three-dimensional images. On the other hand, it's somewhat limited and finicky: images must be exactly the same size, you can't do any editing after the fact, and the plug-in is prone to crashing if it's unhappy with any of the four images. Try it and see if you like it.

## The Artistic Filters

The GIMP's *Artistic* filter menu contains filters to make a photograph look like an oil painting, sketch, woven rug, or other work of art.

It's a great place to explore. Nearly all of the filters there do something fun, and most of them have lots of parameters to tweak. You can see easily what the parameters do, so it's easy to find an effect you like.

In GIMP 2.2, several of these filters lived in the *Script-Fu* ► *Alchemy* menu.

*Apply Canvas* gives a canvas-like texture to the image (Figure 7-16). *Clothify* gives a similar effect, but with a slightly different texture.




Figure 7-16. The *Apply Canvas* filter

*Cartoon* simplifies colors and adds black outlines near edges with the intent of making a photograph look like a hand-drawn and inked cartoon. Most of the time, it doesn't look much like a cartoon, though. Try it and see what you think—it can give interesting looks somewhat similar to *Posterize*.

*Cubism* turns your image into cubist art. *Oilify* blurs it in ways that make it look like an oil painting. *Photocopy* makes the image look as though it was copied on a poor quality copy machine, with burnt-out lights and overdone edges. It's not a pretty effect, but it may be useful for some projects.

*Predator* turns the image into something that should remind you of the movie *Predator*, or a video game from the seventies (Figure 7-17).

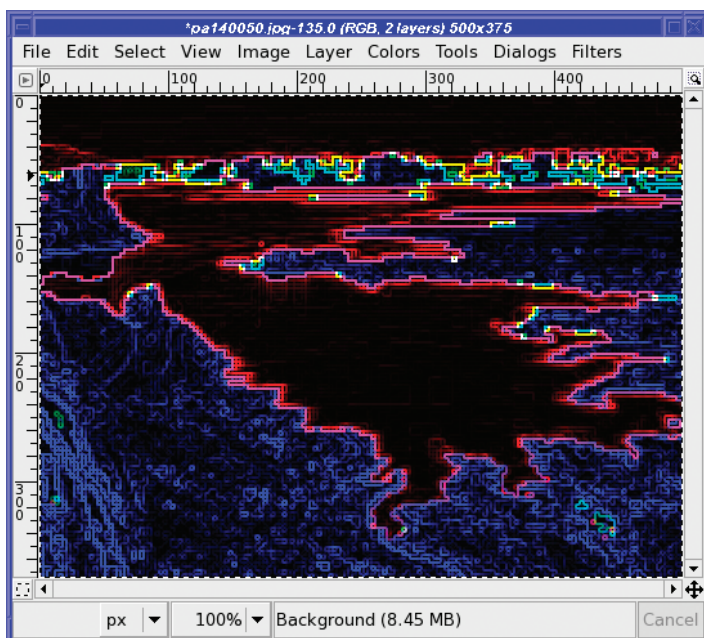


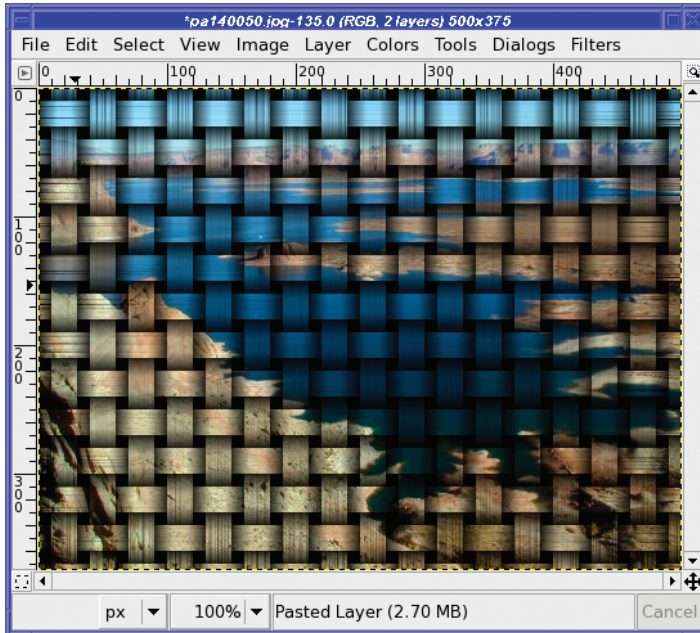
Figure 7-17. Predator effect

*Softglow* gives a bathed-in-light impression, as though the image was hazy and a bit overexposed.

*Van Gogh (LIC)* gives an image a rather subtle shift. Alas, it won't make your photograph look like Van Gogh painted it, but it can create some interesting effects, and it has all sorts of knobs to twiddle. Play with it and see what you can create. *LIC* stands for “Line Integral Convolution,” the mathematical operation that is happening behind the scenes.

*Weave* is another texturing plug-in like *Apply Canvas* and *Clothify*. It gives the nifty effect that the image has been woven into a basket (Figure 7-18).





**Figure 7-18.** *Weave* creates a basket.

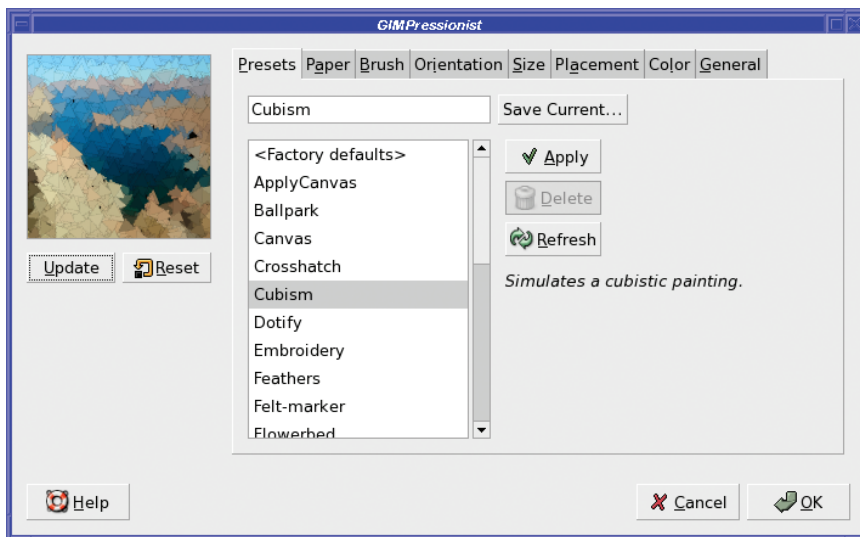
---

**Caution** Weave is unlike most filters: its Undo doesn't work in one step. It takes three Undo operations to get rid of a Weave. Don't be too dismayed when the first Undo doesn't do what you expect; just hit a couple more and you'll be fine. You'll learn in Chapter 11 what this means in terms of how Weave was written—and how to fix it if it bothers you!

---

## GIMPressionist

The crown jewel of the *Artistic* menu is *GIMPressionist*. It has an enormous collection of options and can duplicate the functions of several other plug-ins. The catch is that it's a bit harder to use (Figure 7-19).



**Figure 7-19.** *GIMPpressionist offers many options, such as Cubism.*

GIMPpressionist offers eight tabs, each full of options. *Don't panic!* You can accomplish most tasks simply by using the presets in the first tab, and never looking any farther than that.

When you first try GIMPpressionist, you may wonder why it doesn't seem to be doing anything. You're exploring the presets, but selecting different options seems to have no effect. There are two reasons.

First, GIMPpressionist doesn't have an automatic preview. Whenever you change an option and want to see it reflected in the preview, you need to click *Update* under the preview area.

But even then, changing presets doesn't do anything. If you click *Update*, you just get another version of the preset you used last, not the one you just selected. What's going on?

The answer is that the presets change values in all the other tabs: *Paper*, *Brush*, *Orientation*, *Size*, and so on. Since this would overwrite any changes you may have made in those other tabs—and there's no easy way to undo that—GIMPpressionist requires you to take two steps in order to apply a preset: choose the preset in the list, and then click the *Apply* button to the right of the list. Now, when you click *Update* you'll see the new setup.

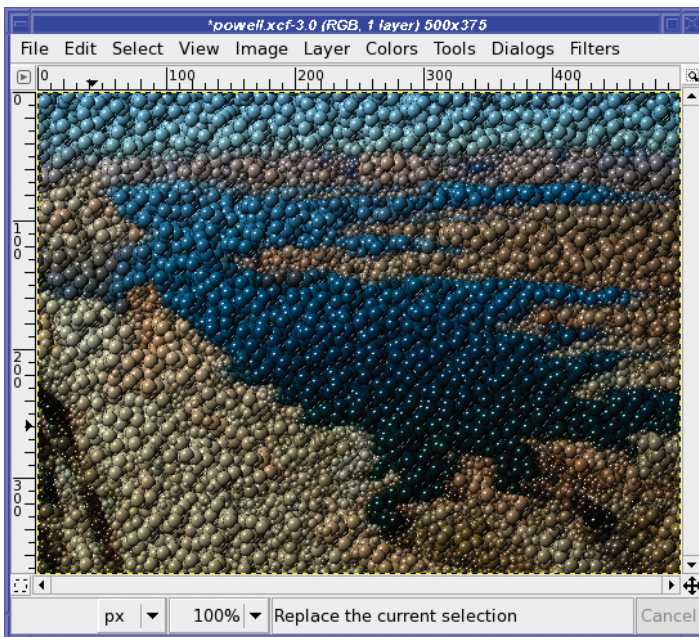
You can spend quite a while just viewing all the presets in GIMPpressionist. It's often hard to tell from the small preview pane exactly what an effect will look like at full size, so you'll often want to click OK, then Undo and bring up GIMPpressionist again (a fast way to do that is *Filters* ► *Re-Show "GIMPpressionist"*, which is bound to Shift+Control+F).

In your explorations, you may find a preset that is almost what you want except for one or two minor attributes. That's your cue to explore the other seven tabs.

If you've selected a preset you like and clicked *Apply*, the other seven tabs will show the settings that create that effect. For instance, the *Apply Canvas* preset uses a *paper* called struc.pgm that has a canvas-like texture; the other tabs are set to reasonable defaults to make the canvas texture look right when it's applied to your image. *Cubism* uses a much finer paper (though still with a canvas texture), but it uses a large graduated brush called grad01.pgm, while *Apply Canvas* used a much smaller brush.

You can tweak the settings as much as you like. You can make small changes: start from *Cubism* and change the paper to create cubist art drawn on bricks, burlap, or marble. Or you can make large changes.

*Orientation* adjusts the angle of asymmetrical brushes. You can make the brush strokes all go in the same direction, or make them change directions randomly within limits you set. *Size* controls the size of the brush: to see a strong example of its effect use the *Marble Madness* preset, click *Apply* then *Update*, and then go to the *Size* tab and make both minimum and maximum size much smaller (Figure 7-20). Click *Update* again to see the effect of your changes.

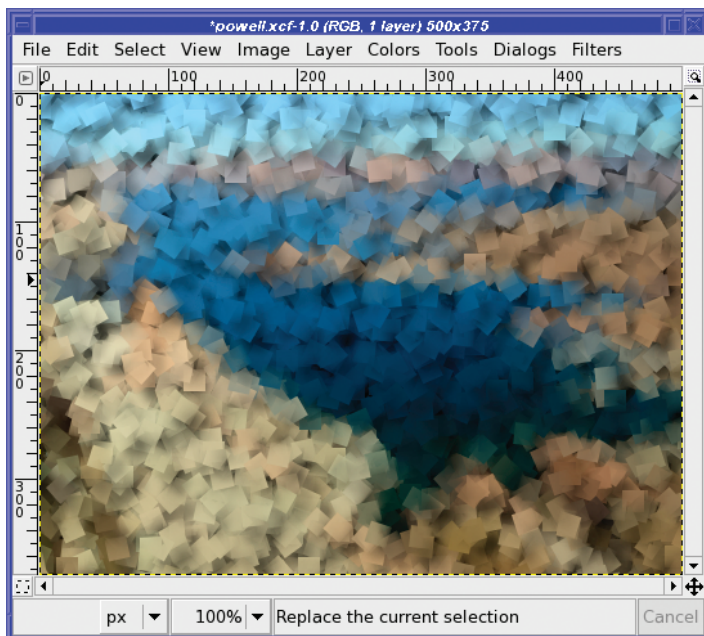


**Figure 7-20.** *Marble Madness* with sizes from 2 to 14

*Placement* controls how the brush strokes will be grouped on the page—whether they’ll be clustered close together or distributed evenly. *Color* controls how precisely the color is laid; for most images it won’t be obvious, but if you’re using a large brush on an image with precise sharp bands of color, this setting might affect how sharp the lines are in the result.

The *General* tab has some useful options. In particular, adding a *Drop shadow* can add depth to an image. Try starting with *Cubism* and then adding a drop shadow: you get an effect like an explosion in a sticky-note factory (Figure 7-21).

You can even make your own papers and brushes. Most of them have extensions of .pnm or .ppm. These files are a special format called PPM or *Portable PixMap*; PGM, or *Portable GrayMap*, is a special grayscale case of PPM. GIMP has no problem reading and writing PPM and PGM files. Try searching for one of the brush or paper images on your system, like blob.ppm, and opening it in GIMP. Make some changes to it and save it as the same type. GIMP will ask you whether you want *Raw* or *ASCII*; use *Raw* (that makes the file smaller, so GIMP can load it faster).



**Figure 7-21.** An explosion in a sticky-note factory

Save it where? Find your GIMP profile (look in the *Folders* category of the Preferences window if you don't remember where it is). Go to the folder called *gimpressionist*. Create folders inside it called *Brushes* or *Papers* if they aren't already there (it doesn't hurt to create both folders if you think you might ever need them). Then save the new brush or paper to the appropriate folder. You'll have to dismiss the GIMPressionist dialog and call it up again to see the new brush, but you don't have to restart GIMP.

You may wonder about the drop-down menu in the *Brush* tab with the *Save As* button next to it. Is that a way to choose a brush, or save one?

The sad answer is no. Choosing a brush in that menu doesn't choose it in GIMPressionist. If you open an existing GIMPressionist brush in GIMP and then try to save it with that button, you'll get an error message because the brush isn't grayscale (though the error message doesn't actually make that clear). If you convert the image to grayscale, then you can save it—but there really isn't any advantage to doing it that way. You're better off ignoring the *Brush* tab's *Select* menu and *Save As* button.

Since GIMPressionist has so many different knobs to tweak, it allows you to save settings when you find a combination you like. *Save Current...* will write the current settings to your GIMP profile, along with a description of what the effect does. Then you can call up the same combination of settings in future GIMP sessions.

The only catch? You probably won't know for sure whether you like a combination of settings until you've seen them in full size on the image. It's hard to tell much from that little preview window. But when you *Re-Show "GIMPressionist"*, the *Presets* tab is back to *<Factory defaults>*. It looks like it's forgotten your changes!

Never fear. Although GIMPpressionist doesn't remember your last preset, it does remember all the other settings. Click *Update* to verify that. If you're happy with your settings, type a name into the field above the preset list, next to *Save Current...*, and then click *Save Current...* GIMP will prompt you for a longer description of your settings.

## The Map Filters

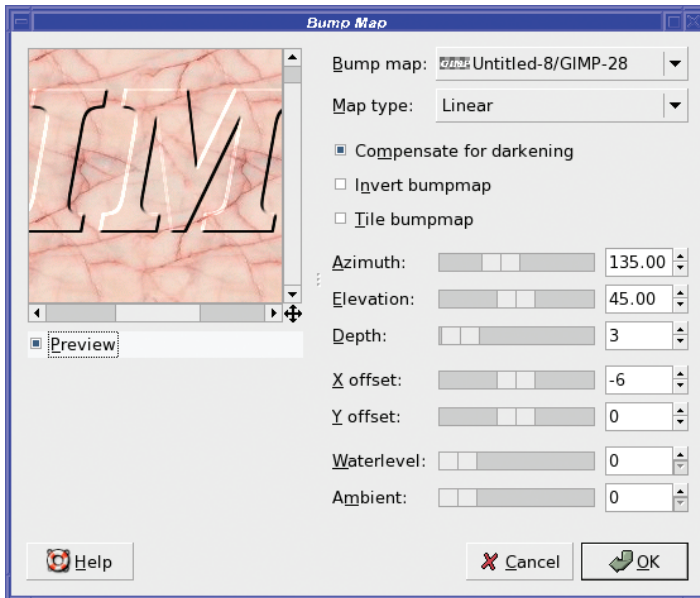
The *Map* filters distort the image a bit, like the *Distorts* filters—but only in very localized and specific ways.

### Bump Map

The first map filter is *Bump Map*, one of the GIMP's most useful filters. It's used in many text and logo effects, but you'll find lots of other applications for it.

It's easiest to see what Bump Map can do if you start with a new blank image and fill it with a pattern (*Edit* ► *Fill with Pattern*). For example, use *Pink Marble*. Then, create a *white* text layer on top of that.

In the Layers dialog, select the background layer, not the text layer. Then call up *Bump Map* (Figure 7-22).



**Figure 7-22.** The Bump Map filter

*Bump Map* makes “bumps” in the active layer, as if another layer (called the *bump map*) was pushing up from underneath.

The map layer is chosen from the drop-down menu at the top of the Bump Map dialog. This menu lists every layer in every image that you have open in the GIMP. It can be confusing: both the image name and the layer name are crammed into the menu. Sometimes there isn't



room for both, so some letters may be left out. Use the tiny layer previews to the left of each image/layer name to help you figure out which layer is which.

Once you've chosen the text layer as the map, you can click *OK* to create the bump map. You'll probably have to make the map layer invisible in order to see the effect. (It's okay if you keep it invisible. You don't have to be able to see a layer in order for Bump Map to use it.)

Bump Map has several options. The *Map type* controls how abrupt the map will be and whether it has sharp edges.

*Compensate for darkening* brightens the image a bit after applying the bump map. The map is implemented using the *Multiply* layer mode (you'll work with modes in Chapters 9 and 10), and that has a tendency to darken the image. Normally, you'll want to leave this option checked to get the brightness back to normal.

*Invert bumpmap* makes the bump appear to go *into* the active layer rather than rising *above* it. If you want an effect like letters chiseled into marble or routed into wood, use *Invert*. (Of course, you can also invert the effect by reversing dark and light in the bump map.)

*Tile bumpmap* makes the bump map repeat across the image. You can use it with a small map image to get a repetitive pattern all the way across a large background area.

*Azimuth* controls the direction of the illumination. The default, 135, makes the light appear to come from the upper-left, but you can light your bump map from any direction. *Elevation* controls how high off the page the light source will seem to be: a higher light source means shorter shadows.

*Depth* controls how tall or deep the bump seems to be.

*X offset* and *Y offset* shift the bump effect in relation to the position of the map layer. Most of the time, you'll already have positioned the bump map right where you want the effect to be, so you'll leave these at 0.

*Waterlevel* controls the "bottom edge" of the effect when your bump map layer has transparency, like a text layer would. Increasing *Waterlevel* can make the effect appear softer and more subtle.

*Ambient* makes the light source appear more diffuse. When it's set at 0, the light will appear to come from only one direction, but increase the value and some of the shaded areas will lighten a bit more to create a more subtle effect.

## Displace

The *Displace* filter shifts pixels in the active layer according to a *gradient map*. The map layer must be the same size as the image you want to displace. (Images that aren't the right size won't appear in Displace's option menu.)

What's a gradient map? It's a monochrome image where the gray value at each point specifies what will happen to the displaced image at that point.

For the Displace filter, horizontal and vertical displacements are handled separately: you can use different gradient maps for the two directions.

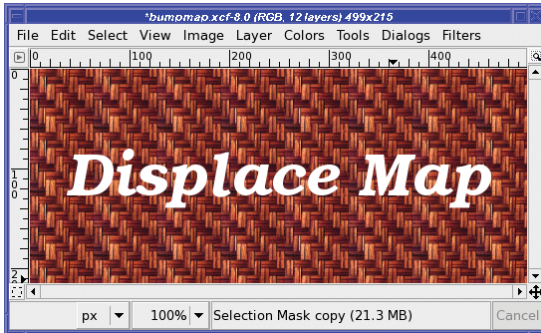
For horizontal mapping, a white pixel in the gradient map means that the corresponding point in the image will be moved to the left. A black pixel means the point will move right, while a medium-gray pixel means the point will stay where it is. Vertical mapping is similar, except white means move up; black, down.

Confused? A couple of examples will help.

Make an image consisting of some text on a patterned background (Figure 7-23).

What sort of map would it take to lay this text out in a diagonal line from upper-left to lower-right?






**Figure 7-23.** Text on a patterned background

That means only vertical displacement, so don't worry about what happens horizontally. On the left side of the image, you want text to move *up*—so that means you want white on the left side of the gradient map. You want the right side to move *down*, so that side needs to be black. A smooth gradient from white on the left to black on the right ought to do that. Make a new layer (call it “gradient”: it's important to give layers clear names to make it easy when you choose them from plug-in menus). Choose the Gradient tool and drag from the right edge of the image to the left to make your gradient.

You can move the gradient layer down in the layer stack, or even hide it. It doesn't have to be visible to work as a displacement map.

Now make the text layer active and call up *Displace...* (Figure 7-24). Uncheck *X displacement* since you won't be needing it, but make sure *Y displacement* is checked. In the menu next to *Y displacement*, select the layer with your gradient. You may not be able to read the full layer name in the menu, but look at the layer previews. You should be able to find your gradient there.



**Figure 7-24.** The *Displace* dialog

Once you've selected the gradient, the preview shows that the text is indeed slanting from upper-left to lower-right, as intended. The background is also slanting, though that's more difficult to see. You can make the effect stronger by changing the displacement value from 20.00 to a larger number, like 50.00.

How would you go about displacing in both X and Y at once? For this, it's best to think about each direction separately.

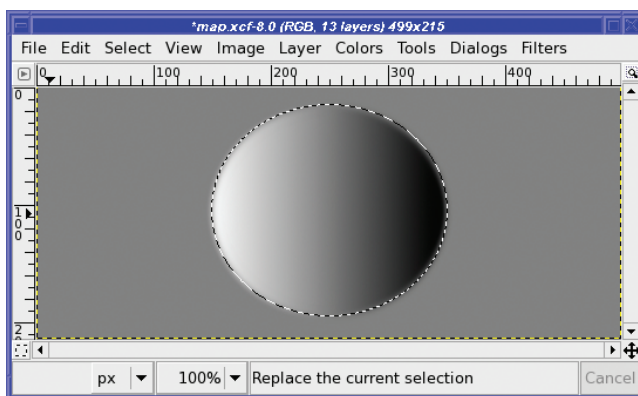
Suppose you want to make it look like you've laid a magnifying glass on the middle of the text. All the text under the glass should be magnified: text near the left edge should move left, text on the right should move right, text above center should move up, below center, down. Everything outside the lens should remain unchanged. How would you design gradient maps to accomplish this?

Consider the horizontal problem first. Make a new layer for your horizontal gradient map; call it "hgrad."

Most of your text will be outside the magnifying glass, so it won't move at all. That means it needs to be medium-gray. Set the foreground color to medium gray (V = 50% in the color chooser), and fill the layer with that color. As long as you have that foreground color handy, this is a good time to make a second layer ("vgrad") and fill that with gray as well. Select the "hgrad" layer again so that you can draw into it.

Now make a circular selection where you want your magnifying glass to be. Give it quite a bit of feathering so the effect will be gradual at the edges.

What should go inside this circular lens? Think about the goal: text at the left edge of the lens should move farther left, so the map needs to be white there. Similarly, the right side of the lens needs to be black. Set the foreground color back to black (the black/white button under the color swatches in the Toolbox are a handy shortcut) and drag the Gradient tool across your selection from right to left (Figure 7-25).



**Figure 7-25.** The horizontal gradient map for a magnifying lens

Now make your vertical lens gradient the same way. It needs to be white at the top, black at the bottom. You can select the "vgrad" layer and re-use the same selection.

Cancel the selection, make the text and background layer active, and call up *Displace...* Choose the "hgrad" and "vgrad" layers for X and Y displacement; you might want to increase the displacement a bit to exaggerate the effect (Figure 7-26).




Figure 7-26. Lens effect using Displace

If you displace in *Polar* mode, the two maps are interpreted somewhat differently. Instead of X and Y, you're mapping *Pinch* and *Whirl*. In *Whirl*, white rotates the pixel clockwise around the center of the image, while black rotates it counter-clockwise. In *Pinch*, white pinches the center of the layer in, while black bows it out. Polar mode is hard to visualize, but if you play with it you might get some good results. Try starting with 0 in the *Pinch* and *Whirl* numerical fields, and then changing just one of them by small amounts to visualize what each one does.

## Some Mapping Toys

*Fractal Trace* turns your image into a *fractal*. This is a class of mathematical equations that produce all sorts of pretty pictures. There are several parameters to play with, whose details are primarily of interest to mathematicians.

*Illusion* creates ghostly copies of parts of your image. There aren't many options to adjust, so give it a whirl.

*Map Object* is another wonderfully useful GIMP workhorse. It takes an image and wraps it around some other shape, such as a sphere or a cylinder (Figure 7-27). It has lots of options, but they're fairly self-explanatory.




Figure 7-27. Map object to a cylinder.

## Tiling

*Make Seamless* warps your image so that the right edge can match up to the left edge, and the top to the bottom. This is useful if you want to *tile*: use several repetitions of a small image as a background for a larger page. It doesn't always give a realistic result, but it's useful anyway. You'll use *Make Seamless* for making backgrounds in Chapter 9.

*Paper Tile's* description is that it “cuts the image, and slides the pieces.” It's not a very realistic effect, though.

*Small Tiles* replaces the image with scaled-down versions of itself, repeated several times ( $2 \times 2$ ,  $3 \times 3$ , and so on).

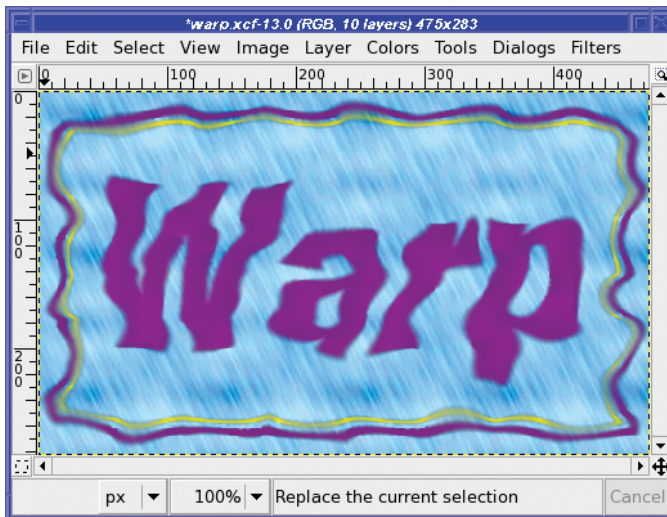
*Tile* is similar but more flexible: it tiles the image at its full size into a larger image whose dimensions you specify. If you have a  $320 \times 240$  image and you tile to a  $640 \times 480$  image, you'll see four copies of the original. You can specify any dimensions, though, not just multiples of the original image size.

If you've just run *Make Seamless* to create a background and want to see how it looks, either *Tile* or *Small Tiles* can help. They're also useful if you're printing many small copies of an image.

## Warp

*Warp* is like *Displace* but with a twist—literally. Like *Displace*, it uses a gradient map to control how much each pixel will move. But *Warp* only needs a single map where *Displace* uses two. When warping, the *difference* between a pixel in the map and the pixels nearby controls the amount and direction of displacement in the source image.

It's most commonly used with random maps. To try it out, create a new layer the same size as the one you want to warp, and then generate some noise (*Filters* ► *Render* ► *Clouds* ► *Solid Noise* works well). Use the noise layer as the warp map (Figure 7-28).

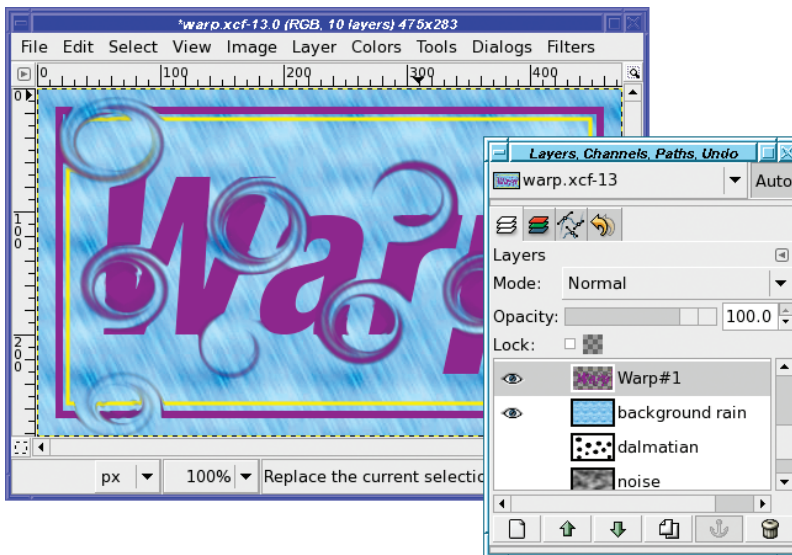


**Figure 7-28.** 25 steps of random *Warp* on some text and a rectangular border. The “Rain” pattern background was not warped.

*Step size* controls how far a pixel can move in each step; *Iterations* controls the number of steps that will be run. Usually, you'll want to run quite a few steps (Figure 7-28 used 25 steps). It's fun to watch the effect of the warp as it proceeds.

In *Advanced Options*, *Dither size* can add some extra randomness. *Rotation angle* is “The angle through which the local gradient of the displacement map (control matrix) will be rotated before displacing the pixel in that direction.” (Well, the dialog *did* warn you that this was an advanced option.) The plug-in's author recommends staying with 90° in most cases. You can add a *magnitude map*, another monochrome image that adjusts the intensity of the effect at each point (white means maximum effect, black gives no effect). There are several even more advanced options: experiment if you're so inclined.

You don't have to use random maps. Try using a series of black feathered polka dots: start with a white background, choose the Ellipse Select tool and turn on feathering, and then make a small circular selection and fill it with black. Make another circular selection and fill it, and repeat until you have as many dots as you want and the layer looks like a dalmatian. Use that as a map for Warp: it's a neat effect and it's fun to watch it run (Figure 7-29).



**Figure 7-29.** *Warping text with a “dalmatian” polka-dot map*

## Adding Patterns to a Layer

The **Filters** ▶ **Render** submenu offers you a collection of patterns you can add to the current layer. None of them are based on what's already in the layer, but they add to it.

### Clouds

*Clouds* provides three ways to get random cloudy patterns. *Add Fog* looks like an overcast sky—except that by default it uses an orange color. (Why orange? I'm sure there was a reason once, but no one remembers it now.) Change it to gray if you want an effect like real clouds, or to any other color you want for other effects.

*Plasma* creates a wild blending of many colors. It doesn't look like clouds, but it makes a nice background.

*Solid Noise* looks like a lumpy surface photographed in black and white. It's useful in effects (you saw it used with the Warp plug-in in the last section) even if it doesn't look that good by itself.

## Nature

*Nature* only provides two filters, *Flames* and *IFS Fractal*. They both use the mathematical concept of *fractals* to create intricate patterns.

*Flame* doesn't actually create flames. It creates arcs and loops like something you might see if you whirled a can of paint around your head. Experiment with the *Colormap* options to make the effect a bit more colorful. The *Edit* button lets you choose different patterns.

*IFS Fractal* is much more general. *IFS* stands for "Iterated Function System," and with this plug-in you can create fractal patterns that repeat in a variety of different ways.

The dialog presents you with three triangles, which you can drag in various ways to affect the preview. Click on a triangle to select it, and then drag it. You can add another triangle by clicking the *New* button, or use fewer by clicking on one and then choosing *Delete*. You can change the shapes, sizes, and orientations of the objects with the *Move*, *Rotate*, and *Stretch* buttons, or you can edit the *Spatial Transformation* numbers. You can turn the triangles into quadrilaterals by adding another triangle and moving it a bit, into pentagons by adding a fifth and moving it, and so forth.

Choose *Full* in the *Color Transformation* tab to add a bit of color to the fractal you create, though controlling it is difficult.

Actually, controlling anything in the IFS Fractal dialog is difficult. The key is to make small changes one at a time. Happily, Control+Z (Undo) works in the dialog, so you can back out changes that didn't work. If you want more general fractal patterns, keep reading: you may like *Fractal Explorer* better.

## Pattern

*Pattern* offers a variety of—well, patterns. *Checkerboard* is obvious (and the *Psychobilly* option can be fun).

*CML Explorer* generates *Coupled Map Lattices*. Mostly, these are patterns that "bleed down" from the top of the image in random and sometimes interesting ways. You can choose from a large variety of functions that will be of interest to mathematicians. Everyone else, try making random choices and watch how the image changes.

*Diffraction Patterns* is a real kick. Diffraction is the way light bends when it passes by an obstacle, or through a slit or a pinhole. This causes the light waves to interfere with each other in interesting ways. The *Diffraction Patterns* plug-in can simulate some of these behaviors, giving you mandala-like images that would look at-home on a sixties wall.

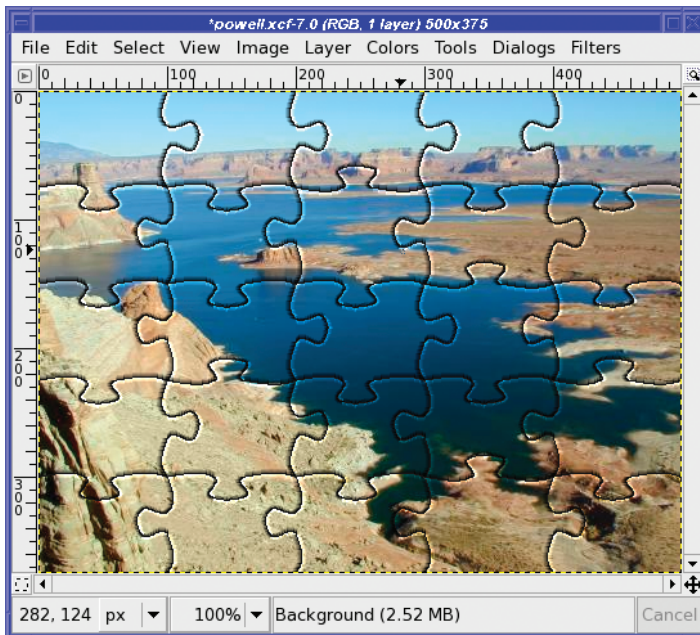
There's no automatic preview: when you make a change you'll have to click the *Preview!* button. Changing colors won't do what you expect: increasing a color won't necessarily show more of that color in the final image. Don't miss the values in the *Other Options* tab, which can change the image quite a bit.

*Divisions* can draw lines on your image to divide it into pieces, but the interface is obscure; for that, *Grid* is much easier to use. You can choose the color of your grid lines, the spacing between the lines, and their offsets from the top and left sides of the image.



You can even make dashed grid lines by specifying a different color for the intersections of the lines. If you set the width of the grid lines to zero, Grid will draw “plus” characters at the intersections of the grid lines without drawing the lines themselves.

*Jigsaw* turns your image into a nifty jigsaw puzzle (Figure 7-30). You can control the number of puzzle pieces, whether they’re square or curved, and how heavily they’ll be shaded.



**Figure 7-30.** *Jigsaw* turns an image into a jigsaw puzzle.

*Maze* generates a random maze in black and white. You can even make one that’s tileable so you can use it as a background for large images.

*Qbist* generates random colorful patterns. The dialog shows nine sample patterns. If you like one of them, click on it: that pattern will center, and GIMP will offer some other patterns similar to the one you chose. If you like one of those patterns better, click on it.

Continue until you have one you like; then click OK, and your layer will fill with a pattern similar to the one you saw previewed.

If you find one you *really* like, click *Save* to save the pattern to a .qbe file. You can *Open* it later if you want to use it again.

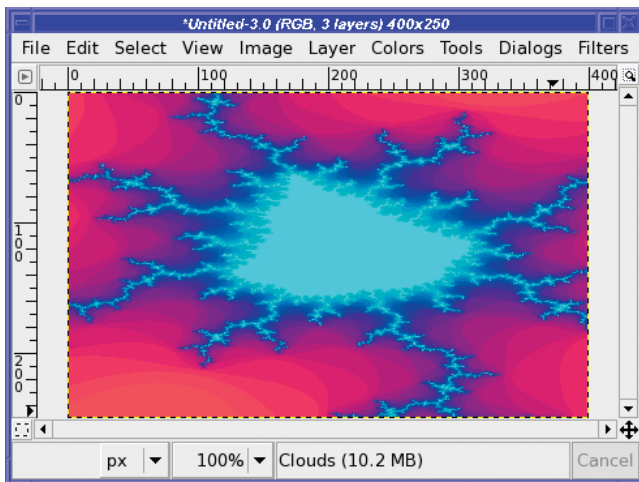
*Sinus* won’t give you a headache—it makes patterns based on a combination of sine waves. It can produce anything from shaded bars to the grain pattern you see in a really nice piece of wood. It uses two colors, which you define in the *Colors* tab.

By default, you’ll merely get straight bars shading into one another. But vary *X scale* and *Y scale* (make sure they’re not the same value), and then play with the *Complexity* to get much more interesting patterns.

*Circuit* is outside the *Pattern* submenu, but it’s similar to *Maze* except that it draws rounded traces in black and white. They look like the patterns on a printed circuit board.

## Fractal Explorer

If you like beautiful fractal patterns like Figure 7-31, you can spend hours in *Fractal Explorer*.



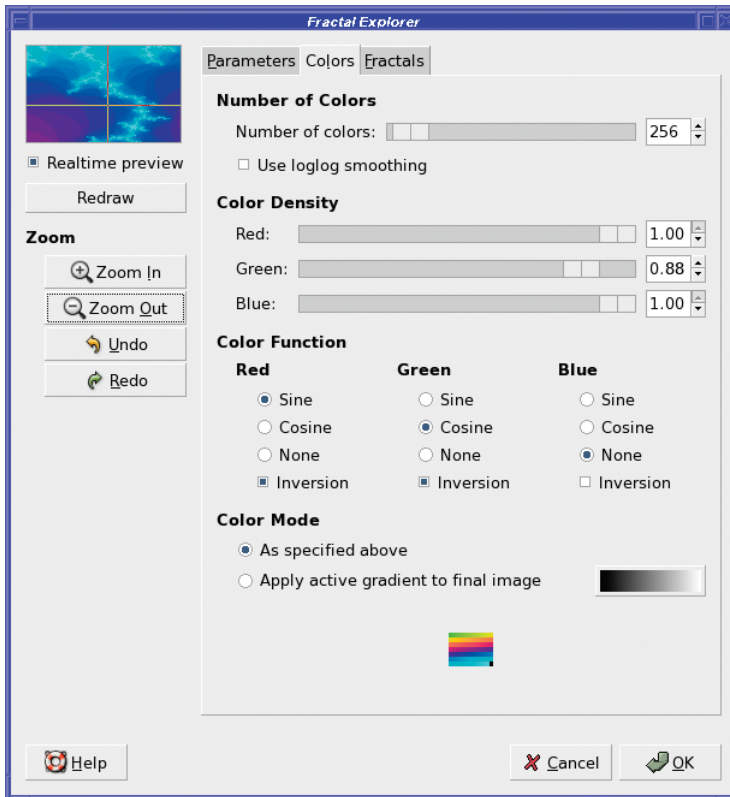
**Figure 7-31.** “Energetic Diamond”—a type of Mandelbrot fractal

If you’ve never played with fractals before, start with the *Fractals* tab. The list there gives you a sampling of patterns you can generate. Double-click on anything in the list to see a preview.

Once you find something you like, you can tweak it as much as you like. Go back to the first tab, *Parameters*. Twiddling any of the parameters will change the shape of the fractal. Make small changes: fractals are very sensitive, so a tiny bit of fudge in a parameter can mean a huge change in the image. The *Colors* tab lets you adjust how much of each color will be used in the image.

One interesting property of fractals: the more you zoom in on a fractal pattern, the more detail you see. The *Zoom* buttons in Fractal Explorer let you home in on the center of a fractal—but that’s not really the point. The most interesting detail may be in one of the fractal’s outer arms, as it is in Figure 7-31. Fortunately, you can also drag in the dialog’s preview area to zoom in on whatever area interests you most (Figure 7-32). If you don’t like what you see, you can *Undo* or *Zoom Out* and try another region.

Since you can spend a lot of time tuning fractal parameters to get exactly the right curve, there are buttons to *Save* a pattern you really like in the *Parameters* tab, or *Open* one you’ve previously saved.



**Figure 7-32.** Zooming in on a region in *Fractal Explorer*


## Gfig

*Gfig* is an odd beast: a vector-drawing program built inside a raster image-editing program.

Its dialog is a simple vector-drawing application all by itself (Figure 7-33). You can make lines, rectangles, ellipses, Bezier curves, stars, spirals, and other patterns. You can move objects around independently and change their fill or stroke style.

What's the catch? Well, *gfig* doesn't save to a standard format. It draws into its own layer in your GIMP image, but if you want to edit the shapes you made later as *gfig* objects, you'll have to save to *gfig*'s own format, using *File* ► *Save* from the *Gfig* dialog.

You probably don't want to use *gfig* for creating elaborate vector-graphics art. But for adding a few figures to a GIMP image that you might want to change or resize later, it can be quite helpful.



**Figure 7-33.** *Gfig* is a vector-drawing application built inside of GIMP.

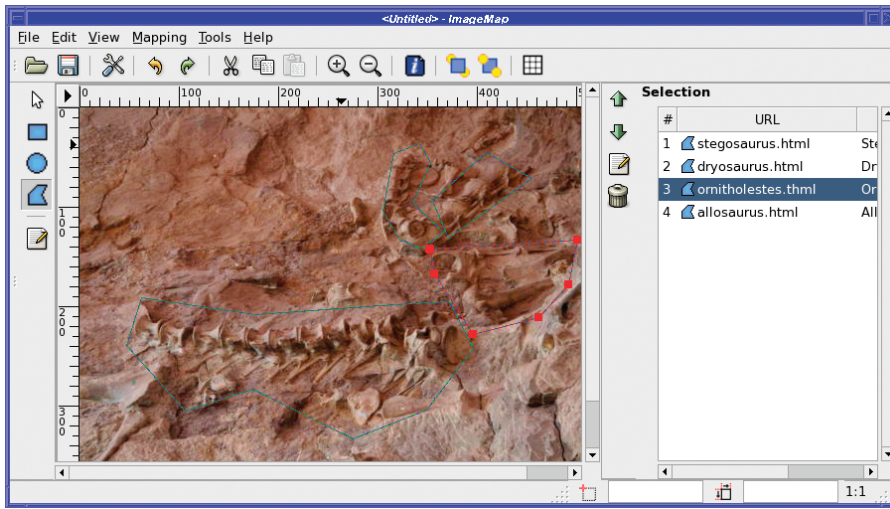
## Other Patterns

- *Lava* creates a red-on-black pattern that looks a bit like glowing lava at night.
- *Line Nova* creates a starburst pattern using the current foreground and background colors.
- *Sphere Designer* creates a sphere using a variety of complicated options. You can control all aspects of how the sphere is rotated and its apparent texture. In GIMP 2.2 and earlier, the options were very difficult to figure out, but they've been simplified quite a bit in 2.4. Play around and see what you can make. (For a simpler sphere plug-in, keep reading: there's one in the Xtns menu of the Toolbox.)
- *Spirogimp* creates a pattern like a “spirograph” toy. You can control the number of teeth each spirograph wheel has, as well as the colors, drawing tool, and pen that will be used.

## Filters to Help Make Web Pages

The GIMP offers a few filters to help you prepare images for your web pages.

*ImageMap* lets you specify regions of the image that will link to other pages. You can specify rectangular or circular shapes, or use the Polygon tool to make arbitrary shapes (Figure 7-34).



**Figure 7-34.** Define regions in ImageMap that will take the user to different web pages.

To define a rectangular or circular region, click at one corner, move to the other corner, and click again. A dialog pops up to let you enter the website for that region.

Polygonal regions are only slightly more complicated. Click on each point in turn, building up the region. When you get to the last point, double-click instead of single-clicking to complete the region and bring up the dialog.

You can select regions and edit or change them by clicking on them in the list on the right side of the dialog.

When you've defined all your regions, you can save the image map using the ImageMap dialog's *File* ► *Save*....

The filters *Py-Slice* and *Perlotine* are two ways of chopping up an image into a table of smaller images to put on a web page. *Py-Slice* is written in Python and *Perlotine* in Perl (so in GIMP versions prior to 2.4, look for them in the menus with those names), but otherwise they're pretty similar.

Drag guides from the left and top rulers in the GIMP image window to define where the image will be chopped.

Either *Py-Slice* or *Perlotine* will chop the image into bits, and then generate a web page in which the bits are reassembled as a table. You can specify the type of image that will be saved (GIF, JPG, or PNG), where the HTML and new images will be written, whether there should be space between the images in the table, and a few other details about how you want the web page to work.

This is useful for lots of different sorts of web pages. You can use an image of a keypad to let the user enter numbers; or you can split a map into quadrants and let the user explore them separately. Or you might just want to display extra space between pieces of your image.

The GIMP won't show you the web page: it doesn't know how to show HTML! You'll have to open the HTML file in a web browser to see the result. Of course, there won't be any links on the individual images yet; you'll have to edit the page to make it do what you want.

If you specified no spacing between table elements, don't be too surprised if the web page looks like the original image. It should! Nothing has been deleted from the image; it has merely been split with a very sharp knife.

The act of splitting up an image following the position of the guides is known as a *guillotine* operation. If you just want to split up one image into several smaller images, but you don't need an HTML table, *Image* ► *Transform* ► *Guillotine* is for you. It will open all the new images in GIMP, and you can take it from there.

The final filter in GIMP's *Web* menu is *Semi-Flatten* (in some versions it can be found in *Filters* ► *Colors*). *Semi-Flatten* provides a solution for browsers such as Internet Explorer that don't support full PNG transparency. *Semi-Flatten* will preserve areas of full transparency, but pixels that are only partially transparent will be given some combination of their current color and the background color. Before you use this plug-in, be sure to set the background color to match the page where the image will go.

## Animation Helpers

The *Filters* ► *Animation* menu offers optimizers to make your animated image more efficient. *Optimize (for GIF)* was discussed in the animation project at the end of Chapter 3: it attempts to reduce GIF file size, and adds layer names that help you control the speed of the animation. *Optimize (Difference)* performs more general optimizations, not specifically targeted at GIF animations. *UnOptimize* removes some of the optimizations that the other filters add; this can simplify layer names.

*Playback* is the way to preview an animation in GIMP.

The *Animators* menu offers a few filters that create animated images for you. Each of these starts from one or more existing layers and creates a new animated image.

*Blend* makes an animation that gradually fades between the layers of the current image.

*Burn-in* takes an image that consists of two layers—a text layer and a background layer—and creates an effect of the text being “burned” into the background, letter by letter.

*Rippling* and *Waves* each make an image look as though it's being viewed through rippling water. *Waves* is the better effect, though you'll want to run it on an image that has only one layer. *Rippling* works on the active layer and ignores any others.

*Spinning Globe* wraps your image around a globe and then rotates the globe. It works best on an image with a single layer: if there are multiple layers, you'll see flashes as the animation switches layers, but only one layer will be animated. It also works best if your image is tileable, or at least if the left and right edges join together smoothly.

## Alpha to Logo

The *Alpha to Logo* scripts (in the *Script-Fu* menu of GIMP versions prior to 2.4) create interesting effects from a text layer (or any other layer with transparency). If you want some text in a fancy style, try some of these scripts! Figure 7-35 shows the collection, though of course they all have options to let you change fonts, colors, and other aspects of the effect.

All these logo scripts are also available in the *Xtns* ► *Logos* menu of the Toolbox (in earlier versions, *Xtns* ► *Script-Fu* ► *Logos*). If you don't already have a layer you want to use, the Toolbox versions will create a brand-new image containing your fancy text.





Figure 7-35. The logo scripts

## BENDING TEXT INTO OTHER SHAPES

One of the logo scripts isn't shown in Figure 7-35 because it does something a bit different: the very useful *Text Circle* bends text in a circle without adding any extra styles to it.

You may remember that *Polar Coords*, from the *Distorts* menu is another way to bend text into a circle, and that *Curve Bend* can change the vertical height of a layer. But what if you want to bend text around a more complex shape?

The answer is a button in the Text tool options: *Text along Path*, which can take a text layer and warp it along any path you've defined with the Paths tool.

Define the path first, and make sure it's selected in the Paths dialog. Then create the text layer. If you're using a text layer you've already made, make it active by choosing the layer, selecting the Text tool in the Toolbox, and then clicking on the text in the image window.

This should enable the *Text along Path* button, which will create a new path consisting of the outline of the text, warped so that each letter follows the old path.

Since the result is a new path, you have several options for what to do with it. You can convert it to a selection and fill it with a color or pattern, edit it further, or export it as an SVG file.

## Decor

The *Decor* menu lets you dress up an image—or dress it down. You can add a bevel around the outside of an image, add a colored border or make the outside of the image look fuzzy, add splotches that look like coffee stains, or make the image look like an old photo you found in a drawer.

You can round the corners, or make it look like a filmstrip or unmounted slide—these two only work on images *without* an alpha channel. If they're grayed out, you'll have to use *Image* ► *Flatten* to use them, which will combine all your layers into one image and eliminate any transparency.

If you have an image that's in grayscale mode, you can carve it into another (colored) image with *Stencil Carve*, or use it as a map to create interesting chrome effects with *Stencil Chrome*.

## Scripts to Make New Images: The Xtns Menu

The Toolbox's *Xtns* menu, like *Filters*, offers a collection of scripts and plug-ins for lots of different effects. The difference is that the Toolbox scripts all make a new image and don't require you to have an image already open.

*Buttons* offers two scripts to make buttons with text on them, which you can use on web pages or in programs. You can probably make more interesting buttons by searching for web tutorials, however.

*Logos* offers the same list of logo styles as in *Filters* ► *Alpha to Logo*, except that you don't need an image open to use them. I usually find it more convenient to make logos from the Toolbox, and then paste them into another image.

*Misc* ► *Sphere* lets you make a sphere. It's a bit like *Filters* ► *Render* ► *Sphere Designer*, but a lot simpler. Set the foreground color and the image's background (alas, Sphere can't do a transparent background—but in Chapter 11 you'll learn how to fix that). Specify the lighting angle and the size of the sphere, and whether or not you want a shadow. Voilà! A nice sphere.


*Patterns* offers a collection of designs you can use as background images or for any other purpose you can think of.

A *Truchet* pattern (named for Dominican priest Sebastien Truchet, who first studied them) is a network formed by combining a series of tiles, each of which has a circular arc across two opposite corners. *Truchet* and *3D Truchet* create an image filled with two variants on a Truchet pattern.

*Camouflage* makes a pattern like you might see on army fatigues. You can change the colors and make camouflage for other environments, of course.

*Flatland* gives a random pattern rather similar to the one in *Filters* ► *Render* ► *Clouds* ► *Solid Noise*.

*Land* and *Render Map* create two different but equally interesting “artificial world map” views. *Land* gives textured green hills against a powder-blue sea (of course, you can change the colors by using a different gradient). *Render Map* looks more like something you might see on a globe of the earth. They look even better if you wrap the resulting image around a sphere by using *Filters* ► *Map* ► *Map Object*—make your own planet (Figure 7-36).



**Figure 7-36.** Make your own planet with *Render Map* (left) or *Land* (right), combined with *Map Object*.

---

**Tip** You might want to take a look at the layer name after you run *Render Map*.

---

*Swirl-Tile* makes an interesting swirly pattern. You can tile it to a larger image to make a place mat.

*Swirly* gives a pattern of black-and-white squares with a swirl in one of the corners.

Finally, *Web Page Themes* offers several “looks” you might want to use on a website. Under each theme, you can make several different types of element, such as arrows, buttons, or headings, all coordinated. If you happen to like one of the motifs, you might find this useful.

## Summary

Well, that’s quite a list of toys! When you first installed GIMP you probably played around with a few of the filters, but it’s easy to miss some of the most useful functions GIMP has stuffed in its menus.

Now you should know when to look in the Toolbox menus versus the menus in the image window. You’ve experimented with the GIMP’s offerings for changing lighting and shadow, adding noise, and detecting edges. You can map images to three-dimensional objects with the Map filters and turn them into art with the Artistic filters. You know how to create interesting logos and other text effects, and how to make images for your web pages.

Some of these techniques are helpful for working with photographs, but you’ll find that they’re even more useful when making drawings. You worked with basic drawing tools in Chapter 4, but now that you’re familiar with heavy-duty image-processing techniques, you can go far beyond the simple examples from that chapter.

So take a firm grip on your mouse (or drawing tablet, if you have one) as we move into the next chapter and learn some techniques for advanced drawing.





# Color Manipulation, Channels, and Layer Modes

In earlier chapters, you learned how to do simple color correction: how RGB and indexed color differ, and how to change the foreground and background colors in the Toolbox. But a deeper understanding of the GIMP's color model can illuminate all sorts of useful image-processing tricks.

In this chapter, I'll describe the way the GIMP represents colors, different color models, and techniques to manipulate colors.

Understanding color requires a more theoretical approach than you've seen so far. A familiarity with color implementation in digital images is helpful to understanding the sorts of advanced tricks that can be played.

If you feel overwhelmed by the theory, you can skip those sections and use this chapter for reference when you have questions about color. Later chapters will not require an understanding of color theory, though some of the techniques covered here will help in other projects.

This chapter will cover the following:

- A review of RGB and CMY color
- Working in HSV
- Working for print: CMYK
- GIMP's alternative color choosers
- Working with grayscale or black and white
- Coloring monochrome images and making sepia photos
- Using the Threshold tool to clean up black-and-white scans
- Indexed color
- Picking colors from the image
- The color channels
- How to use color decomposition
- Color profiles

## A Review of RGB and CMY Color

You've already worked with RGB images, and you know the letters stand for *red*, *green*, and *blue*. But do you know why those three colors are used? Why not red, yellow, and blue, like the paints you probably mixed as a child?

The answer has two parts. The first part has to do with the structure of the human eye. Our color vision is made possible by light-sensitive cells in the retina called *cones*. There are three types of cones: some are sensitive to red, some to green, and some to blue. Every color we can see, we perceive as a combination of these three *primary* colors.

The second part of the answer lies with how colors of transmitted light combine to make other colors. Transmitted colors are *additive*. The light of one color, combined with the light of another color, is perceived as a third color that is lighter than either component.


### Additive Colors

Computer monitors use the additive technique to depict color. Each pixel in a computer monitor is a combination of red, green, and blue lights. These three colors combine efficiently into virtually any color that the eye can perceive.

If you hold a magnifying glass up to your screen, you should be able to see the separate colors. It's easiest to see in white areas—that's where all three colors are lit up.

Figure 8-1 gives you an idea of how the primary colors combine to make other colors.

Notice that the center of the circle, where all three colors combine, is white. (It may not look white to you, an optical illusion created by the other colors nearby. Try covering the other colors so that you see only the center portion. You'll see it really is white.)



**Figure 8-1.** The primary (outer circles) and secondary (overlapping areas) additive colors



## Subtractive Colors

So what's up with those red, yellow, and blue paints you mixed as a child?

The red, green, and blue *additive primary* colors work for transmitted light. But ink or paint on paper follow different rules. As you add colors of transmitted light, the image gets brighter. But as you add shades of paint to a reflecting surface, the image gets darker. That's *subtractive* color.

The paint absorbs light of particular frequencies. The light that doesn't get absorbed is reflected to your eye. In other words, the paint subtracts colors from the white light shining on it. When you mix several colors of paint together, the paint absorbs more and more colors, subtracting from the colors that reach your eye. Reflected color is called *subtractive*.

Each of the *additive* colors that come from your RGB monitor has an opposite, or complementary color. The complement of blue is yellow, of green is magenta, and of red is cyan. You can see the relationships in Figure 8-1.

For subtractive colors, it works best to use those complementary colors as primaries: so the *subtractive primary colors* are cyan, magenta, and yellow, and the color model based around them is known as CMY.

Take some yellow paint and put it on white paper. The white surface reflects the full spectrum of magenta, cyan, and yellow light, while the yellow surface absorbs the magenta and cyan—what's left is yellow.

But wait! When you shine differently colored lights on a white screen, you're looking at reflected light, but it follows the "additive" rules. What's up with that?

That's confusing, but the reflection is a red herring. The color you see is still coming from transmitted light, even after it's reflected. It's not subtractive color because there's no tinted surface subtracting from the colors you see.

## The Relationship Between Additive and Subtractive Systems

It's convenient to consider these two systems, additive and subtractive, as mirrors of each other. They have their own rules about how colors work, but the rules are exact opposites of each other.

For example, if you mix any two of the additive primaries (red, green, or blue), you get one of the subtractive primaries (also called the additive secondaries): magenta, cyan, or yellow, as you saw in Figure 8-1.

Conversely, mixing any two subtractive primaries of paint will yield the additive primaries: yellow + magenta = red, yellow + cyan = green, and cyan + magenta = blue. So ultimately, the object of a subtractive approach is to reflect red, green, and blue—aiming the RGB colors of light at your eye, just like your monitor does!

---

**Tip** The GIMP can help you experiment with combining colors, using the *Addition* and *Subtract* layer modes and the color picker. That's how Figure 8-1 was made. You'll learn how in the next chapter.

---

There's another way the two approaches mirror each other. If you keep adding transmitted light colors, eventually you'll get white. If you keep mixing colors of paint (subtracting reflected light), you theoretically end up with black. (In practice, with paint you get a muddy dark brown most of the time, but that's just reality refusing to conform to theory. You'll learn more about black in the discussion of CMYK color, and in the section "Printing" in Chapter 12.)


## Color Depth

Digital images have another property: *depth*, or how many different shades of each color the video system can display. You will sometimes see color depth abbreviated as *bpp*, *bits per pixel*.

You may hear the phrase *bits per channel*, the number of bits used per color in each pixel. Multiply by three (for the three primary colors) to get bits per pixel—or sometimes by four if the image has transparency.

Many early color computers could display only 256 colors—this is  $2^8$ , the largest number that can be expressed in eight bits, and hence this was sometimes referred to as *eight-bit color*.

Later, computers were introduced that could display 16-bit color, or 65,536 ( $2^{16}$ ) colors. Today, most computers can display *24-bit color*: eight bits per color channel, or 16,777,216 different colors. This is sometimes called *true color*. Figure 8-2 compares these three color depths.



**Figure 8-2.** 24-bit, 8-bit, and 16-bit color

Modern computers most often use 32 bits to represent a pixel. Generally, 32-bit color has the same number of colors as 24-bit color. The extra eight bits may be used to represent transparency (as the GIMP does), or it may not be used for anything at all.

---

**Note** Macintosh computers show the color depth choices as “256”, “Thousands”, or “Millions”, corresponding to 8-bit, 16-bit, and 24- or 32-bit color.

---

Some professional graphic artists and motion picture animators use even greater color depth: 16 or more bits per color channel. Normal computer monitors can't display this many colors, nor can normal printers. But extreme color precision can offer more control to the artist who must produce exactly the right color on professional output equipment. Some scanners can also produce images with 16 bits per channel, though you'll never see the difference on a standard computer monitor.

---

**Caution** When someone says “8-bit color,” they might mean 8 bits per pixel—only 256 colors total—or they could mean 8 bits per channel, or 24-bit true color. The same is true of someone saying “16-bit color”: it could mean slightly less than true color, or it could mean high-depth color with 16 bits per channel. Usually you can tell from the context, but the terminology can be quite confusing.

---

The 24-bit model is pretty easy to understand. Divide 24 by 3, and you get 8 bits to assign to each color. However, 16-bit color is not so easy to figure; you can't divide 16 by 3 so conveniently. What gives?

In 16-bit color, the usual technique is to assign five bits each to red and blue, with the remaining six going to green. It turns out our eyes are more sensitive to green light, so distinguishing fine differences of intensity is most useful in that part of the spectrum.

## The Indexed Palette

8-bit color is even weirder. Because there just aren't enough bits to do a good job, there never was a truly common standard bit assignment. The usual trick was to create a *palette* of colors, assigning the values however the designer saw fit, and rounding the input color to the nearest member of the palette for output to the display.


Since monitor phosphors could show far more information than the system could supply, the palette could be redesigned with different goals in mind: making text and windowing systems look good, making computer graphics such as logos look good, making photos look good, and so on. Each window on a display could use a different palette.

Fortunately, you probably won't have to deal with palette headaches any more except for indexed images (see the section on “Indexed Color” later in this chapter).

## Representing RGB Color in Web Pages

Knowing RGB notation can help in understanding color in web pages written in hypertext markup language, or HTML. Although HTML lets you specify colors by name, such as “red”, many web page authors prefer something more flexible. Instead, they specify colors by their RGB values.

Figure 8-3 shows what the color notation looks like. HTML notation shows six digits representing three two-digit numbers in base sixteen (*hexadecimal*), ranging from 00 (darkest) through 99 to ff (brightest). Bright red is shown: the first two digits are *ff*, meaning as much red as possible. The second and third pair are both *00*, meaning no green and no blue, only red.



**Figure 8-3.** HTML notation for “red”

In HTML, you add “#” in front of the number GIMP gives you. So red would be #ff0000 in HTML notation. (You can also write this in a more compact three-digit form, as #f00.)

Notice that you can also read the R, G, and B values off the sliders in GIMP’s color chooser. R is 255 (ff is hexadecimal for 255), while G and B are both 0. If you don’t know how to convert between hexadecimal and decimal numbers, just use GIMP’s color chooser to pick a color, copy the number in the *HTML Notation* field, and then put a “#” in front of it.

## Working in HSV

RGB is straightforward, and it’s the way the computer represents colors internally. But sometimes it works better to represent colors in some other way.

HSV, or *Hue, Saturation, and Value*, is a very useful color model. When you experimented with color correction in Chapter 6 using the Hue-Saturation dialog, you used HSV. You’ve seen at least some aspects of HSV if you’ve used the GIMP’s color chooser.

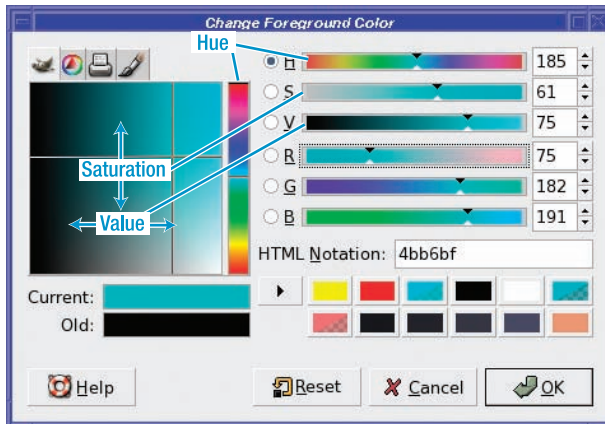
The HSV model, like RGB and CMY, uses three numbers to represent each color. But instead of splitting an image into separate color components, HSV uses conceptual properties of the image.

*Hue* is the color of the image, represented as a single number. In GIMP’s color chooser dialog, it goes from 0 to 360, not 0 to 255 the way the RGB values do. The number represents degrees on a circle: a hue of 360 is the same as a hue of 0 (both red). You can see the hue values laid out in the vertical bar in GIMP’s color chooser, from reds at the bottom (hue=0) through orange, yellow, green, blue, magenta, and finally back to red again (hue=360) at the top.

*Saturation* is how intense a color is. It ranges from 0 to 100. A color that’s bright and vivid is *highly saturated* and has saturation of 100. A color that’s pale and looks “washed out” will have a lower saturation, down to white at the bottom (saturation=0). In GIMP’s color chooser, dragging up and down in the color square changes saturation; or you can drag the “S” slider by itself.

*Value* represents how bright a pixel is, from 0 to 100. A black pixel has value=0; a bright pixel (whether it's white or some other color) has value=100. Dragging left and right in GIMP's color square changes value, as does dragging the “V” slider.

Figure 8-4 illustrates these relationships.



**Figure 8-4.** Hue, saturation, and value in the color chooser

## Why Use HSV?

Why learn this model? Why not just stick with good old Red, Green, and Blue?

Experimenting with the color chooser will provide the answer. Try picking a color using only the R, G, and B sliders. Then try again using the H, S, and V sliders.

Most people find that it's more intuitive to choose colors according to questions like “What color should it be?” and “How bright do you want it?” than it is to figure out how much red you need to mix with how much green to make the shade of yellow you're after.

In addition, other aspects of the HSV model can provide a shortcut to certain image-editing operations.

You've already seen the Hue-Saturation dialog in action: first in Chapter 2, where reducing saturation provided a quick fix for red-eye, and then in Chapter 6, where changing hue helped to get rid of a color error in a photograph. You'll learn other uses for HSV later in this chapter.

## Working for Print: CMYK

Desktop printers use the subtractive colors—with a twist. In order to represent dark colors in CMY, you would normally combine all three inks in fairly equal amounts. But due to the physical properties of ink, this often ends up as a muddy dark brown instead of a nice sharp black. It also uses up a lot of expensive colored ink. So most printers add a fourth color, pure black, and add that to the CMY mixture to get better dark colors.

This might have been called CMYB—but B is already used for blue when talking about colors. So, instead, black is represented by K, and the printing color space is known as CMYK.

**Note** You might also sometimes see it written as CYMK.

---

## CMYK for Professional Printers

People working primarily for print media, on professional equipment, want a lot of control over how much black gets mixed in when an image is printed. Therefore, they like to work directly with CMYK colors when they're editing images.

The GIMP is limited in that respect. The Decompose function, which I'll talk about later in this chapter, has a CMYK option. The color chooser offers a CMYK tab, and the new "Gutenprint" print plug-in (see the section "Printing" in Chapter 12) allows adjustment of K. However, most tools and filters don't give you control over CMYK color, and most plug-ins can't yet save to CMYK. (This is planned for a future version of the GIMP.)

But fortunately, most of us don't need this functionality. You can't see any difference at all on a computer monitor. Desktop computers and home inkjet and laser printers can print perfectly well from an RGB image.

If you take your images to a professional for printing, you can use external programs to convert from RGB to CMYK. There's a GIMP plug-in that uses a program called *littlecms* to convert to CMYK. Or you can just let the printing house do it for you—they know the characteristics of their printers better than you do anyway.

## GIMP's Alternative Color Choosers

Up to now, I haven't mentioned the tabs at the top of GIMP's color chooser dialog. What are they, anyway?

The tabs provide alternative ways of changing color, in case you don't like the standard Saturation-Value square with the Hue bar next to it. The number of tabs and their order may vary from system to system. Some systems may not show any tabs, if the extra color-selector modules weren't installed (they're currently missing from the Mac install). Other systems may have more than four, if additional modules have been downloaded.

The default color chooser is the one you've already used—the standard GIMP color chooser (which has a picture of Wilber on its tab).

---

**Tip** You can right-click on the tabs to see a list of the options available.


---



## The Triangle Color Selector

The “triangle” color selector provides an alternate HSV interface (Figure 8-5). Instead of a vertical slider, hue is represented by a circle (remember from the discussion of HSV how the Hue slider goes from 0 to 360, representing degrees around a circle?). The two small rings are grab handles that you can drag: one drags the point of the triangle around the circle to choose a hue, and the other drags into the triangle to choose saturation and value.

Notice that you still have the familiar set of sliders for H, S, V, R, G, and B on the right: only the left section of the dialog has changed.



**Figure 8-5.** An alternative set of HSV controls

Once you’ve chosen the hue, drag the other white ring back into the triangle to choose saturation and value. Dragging toward the black corner reduces value; toward the white corner, saturation.

The interface in this tab is pretty. But it’s sometimes a little hard to balance saturation and value—they don’t follow equal scales (for instance, the point in Figure 8-5 has saturation and value both set to the halfway point of 50, but clearly it’s not halfway between the dark and light corners of the triangle).

The interface is also a lot slower than the standard GIMP selector—on my machine I find that it can’t always update as fast as I can drag the circle around. But if you have a fast system, you may prefer this color selector.


## The CMYK Tab

The CMYK tab, with a picture of a printer, offers controls to set CMYK colors. If you’re working for a print medium, and know exactly the CMYK hue you’re after, you can choose it here. But beware: GIMP does not use CMYK color internally, so the value you choose in this tab will be converted to RGB for most GIMP operations.

The *Black Pullout* parameter lets you control how much the other colors can be reduced in order to use black instead. The maximum, 100, means use as much black as possible. Adjusting *Black Pullout* automatically adjusts the other colors (but not the reverse).

## The Watercolor Selector

Finally, the trickiest of the selectors: Watercolor (Figure 8-6). When you drag the mouse in the square palette area, the color you drag gets “mixed with” the current color. At first, the color will be fairly pale; drag in small circles within the same color area to make the color gradually darker. Be careful not to slop over into other color zones, unless you want to mix in a different color.



**Figure 8-6.** *The Watercolor selector*

The slider to the right of the palette controls how fast color is added. To be cautious, or to paint with pastel colors, slide it farther down, as in Figure 8-6. If you slide it down all the way to the bottom, though, no color at all will be added.

---

**Caution** The color chooser is one of the few areas in GIMP where there’s no Undo. So if you accidentally add the wrong color, you usually have to start over. You can start over by clicking one of the color presets to the right of the *Current* and *Old* color swatches; happily, white is usually one of the colors shown there.

---


If you’ve done a lot of watercolor painting, this selector may seem very familiar. If you’ve done a little watercolor painting, and when you mix colors, you usually end up with a muddy brown and have to start over—well, this interface may seem familiar!

To try it out, start with white and add color gradually to see what happens. It takes practice and parsimonious mouse use to get good color combinations without ending up at dark gray or brown.

## Working with Grayscale or Black and White

How do you turn a color photograph, such as Figure 8-7, into a black-and-white one?

It turns out there are lots of ways. They all produce similar results, but they’re not exactly the same. There are quite a few ways to map the three color values into a single brightness value. If you want a very specific effect, you may want to try several methods and compare the results. But why are there so many methods?



**Figure 8-7.** How would this look as a black-and-white photograph?

## Methods of Measuring Brightness

What are all the brightness values, and how do they differ?

Since each pixel in an image has a red, green, and blue value, these three numbers must be combined in some way to turn the pixel into a single gray value. The obvious way is to take the *average* of the three:  $(R + G + B)/3$ .

Another simple method is to extract the value of the most prominent color in the image:  $\max(R, G, B)$ . This is *Value*, the *V* of *HSV* color.

But those methods don't always give you a result that matches what the eye sees. The eye is most sensitive to green, less so to the other colors. *Luminosity* (also called *luminance*) is a measure intended to match the way most people's eyes see colors on a computer display. Instead of just averaging red, green, and blue, luminosity is a weighted average that gives the most weight to green, and the least to blue.

In some cases, though, a slightly more complicated method can give better results. To find the *lightness* of a pixel, GIMP takes the value of the color that's most prevalent, and the value of the color that's least prevalent, then averages those two numbers. This usually gives a "flatter" result, with less dynamic range: the brights aren't as bright, and the darks aren't as dark.

How do you create all these grayscale images?

## Grayscale Mode

The most straightforward way to convert to grayscale is to change the image's mode from RGB (24 bits) to grayscale (8 bits) using *Image* ► *Mode*. Without asking any additional questions, this gives you a black-and-white version of your image.

Grayscale mode uses luminance. Often, this is all you need. It's the only method that converts the entire image; the others will convert only the current layer. It is also the only method that will save the image in grayscale mode (if the format you're using, such as PNG, allows such a concept). This can make a smaller image on disk.

Once you tell GIMP an image should be grayscale mode, it won't let you add any colors to it. If you want to add colors later (even in another layer), you need to convert the mode back to RGB (which won't add the colors back; it will merely change the image's mode so that you *can* add colors).

## Desaturate

The second method is the Desaturate function found in the *Colors* menu (in earlier versions, the *Colors* menu is found inside the *Layer* menu). In GIMP 2.4, this gives you a dialog offering a choice between *Lightness*, *Luminosity*, or *Average*. In earlier GIMPs, Desaturate used lightness, without offering the other two choices.

## Hue-Saturation

You may have noticed that, ironically, Desaturate doesn't actually reduce saturation (if it did, you'd be left with *value*, not *lightness*).

If value is what you want, use the Hue-Saturation dialog from the *Colors* menu, and slide the Saturation bar all the way to the left.

Figure 8-8 shows other uses for the Saturation bar. By reducing saturation in an image, you can make a pale-hued image that can make a nice background for a poster or greeting card. Increasing saturation creates a bright, garish image, which again could be useful for posters or other artwork.

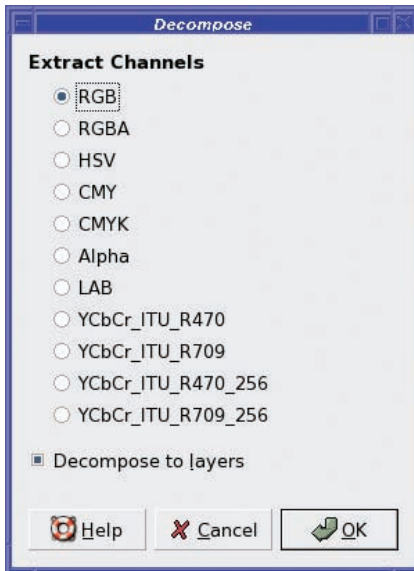


**Figure 8-8.** Reducing (center) or enhancing (right) saturation

## Decompose

No, this isn't what happens when you leave your disk drive out in the rain too long! *Colors* ► *Components* ► *Decompose* is actually a cool way of seeing the color components of your image. (In GIMP 2.0 and 2.2, you'll generally find this under both *Image* ► *Mode* ► *Decompose* and *Filters* ► *Colors* ► *Decompose*.)

The Decompose dialog gives you lots of choices (Figure 8-9). Yikes!



**Figure 8-9.** *The Decompose dialog*

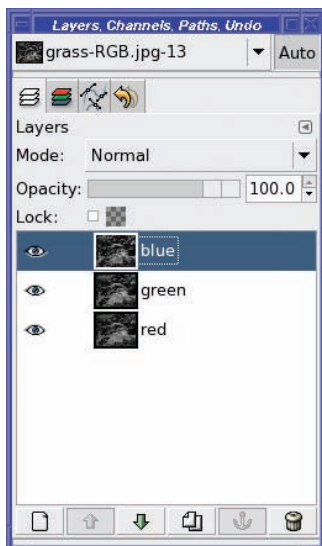
The first five should already be familiar to you. RGB will give you the red, green, and blue channels of your image; RGBA will do the same, but will include information about the alpha channel if your image has transparency. HSV, CMY, and CMYK are all terms you've read about earlier in this chapter. Alpha gives you just the transparency and nothing else.

LAB is yet another color model. The L stands for luminance; A and B represent a color channel between red and green and another color channel between blue and yellow. Some images made in Photoshop and certain other programs use this model. If you don't need to convert LAB images, you probably don't need to bother with it.

The YCbCr\_ITU entries represent specialized color models used for digital video and for the internals of jpeg images. They're similar to LAB, but most people won't find much use for them.

Whichever color model you choose, Decompose will create a new image in which the layers correspond to the color channels (Figure 8-10). If you uncheck *Decompose to Layers* in the dialog, it will give you separate new images, instead of separate layers in a single new image. Either way, it will leave the original image unchanged.

In the new image, what you see is the top layer—in this case, the blue layer. You can view the other layers by clicking on the visibility buttons in the Layers dialog.



**Figure 8-10.** Layers dialog after decomposing to RGB

---

**Tip** You can make changes to these layers, and then recombine them into a color image using **Recompose** (located right next to **Decompose** in the menus). You can also build a color image out of several unrelated grayscale images using **Compose**. This will let you choose among all of the grayscale layers (and only grayscale layers) currently open in the GIMP.

---

Decompose is fun because it lets you see the components of your image in several different color models. But it's also useful: you can use one of the layers as your final monochrome image, if you like it.

Often, the green layer of an RGB decomposition makes a good approximation to what the eye sees. Or try the value layer of an HSV decomposition (which should be equivalent to using the Hue-Saturation dialog and reducing saturation to zero).

The saturation layer of an HSV decomposition is particularly interesting—it looks a little like a black-and-white negative of the original image. The K layer of a CMYK decomposition also looks like a negative. But they're not really equivalent to a negative: if you want to see what a black-and-white negative of the image would actually look like, you can get that by converting to monochrome using any of the methods already discussed, and then selecting **Colors** > **Invert**.



## Channel Mixer

Finally, if you want complete control, try *Colors* ► *Channel Mixer* (Figure 8-11). Clicking on *Monochrome* gives you an opportunity to convert to grayscale using a weighted average like luminosity uses, but with complete control over the weights for each color. Setting red to 21%, green to 72%, and blue to 7% matches the standard luminosity in GIMP 2.4 (earlier GIMPs used 30%, 59%, and 11%). But if your image might benefit from a little more red and less blue, or any other combination, Channel Mixer is for you.

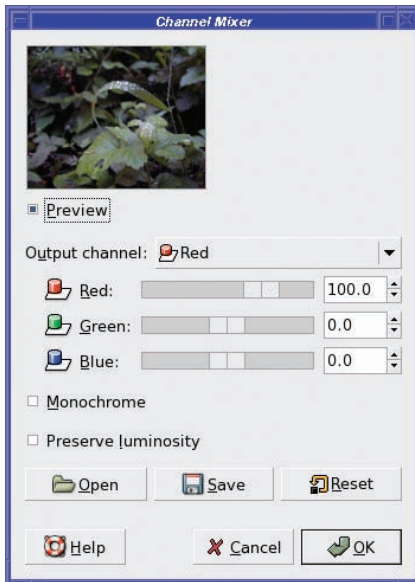



Figure 8-11. The Channel Mixer

Notice that you can make the image much brighter or darker by using the sliders in the Channel Mixer. If you just want to adjust the weights of the colors, without changing overall brightness, the *Preserve luminosity* checkbox will adjust all the values so that the image's luminosity remains unchanged even if the three slider values add to more or less than 100%.

## Coloring Monochrome Images and Making Sepia Photos

Some cameras have a *sepia* setting, which makes a black-and-white photo and then colors it brownish, like a photographic print that's been “toned” to make it last longer.

But if you like the look of sepia prints and your camera doesn't do them, you can make your own with GIMP. Start with a color photo (Figure 8-12).



**Figure 8-12.** A color photograph to be converted to sepia

## Automatic Conversion with the “Old Photo” Filter

First, let's do it the easy way. GIMP already has a plug-in to do this conversion. It's in *Filters* ► *Decor* ► *Old Photo...* (Figure 8-13). Earlier versions of GIMP located this feature in *Script-Fu* ► *Decor* ► *Old Photo*.

- *Defocus* makes the focus a little softer.
- *Border size* creates a fuzzy border around the image, making it look like the edges have been torn. Change this value to 0 if you don't want a border.
- *Sepia* converts the image to black and white, and then tints it sepia (without this, the image will remain colored).
- *Mottle* creates mottling all over the photo, making it look like it's been sitting in a dusty drawer for fifty years.
- *Work on copy* creates a new image, leaving your original unchanged.

By default, it will *Defocus*, add a 20 pixel border, tint sepia, add no mottling, and create a copy.

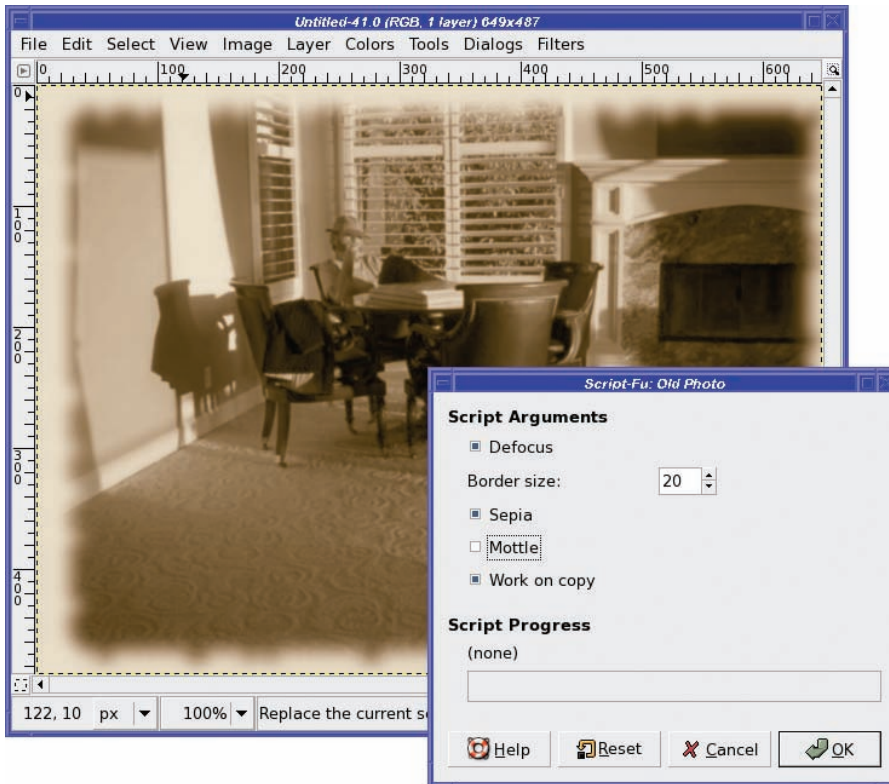


Figure 8-13. The Old Photo plug-in

## Manual Conversion for Fine Control

How does it work? How can you use the GIMP's color system to make your own sepia prints?

Let's try several techniques. They all begin by converting the color image to black and white. But make sure the image ends up in RGB mode, not grayscale, so you'll be able to add the sepia color.

### Colorify

One method is by using the *Colorify* plug-in, accessed through *Colors* ► *Colorify* (Figure 8-14). Colorify tints an entire image as though you were looking through a filter. It starts with a small set of predefined colors, but by clicking on the *Custom Color* button, you can tint with any color you like.

Colorify tends to make images that are fairly dark. You have to make the tint color very light or you may end up with a muddy result.

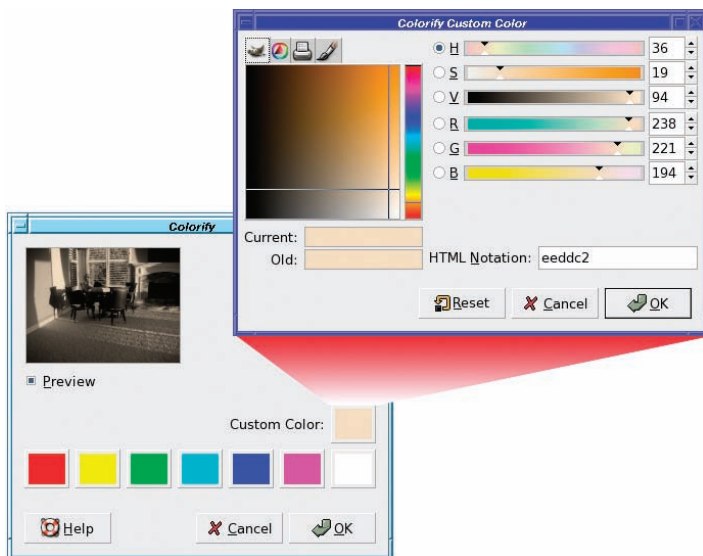


Figure 8-14. The Colorify plug-in

## Colorize

Easier to use is *Colors* ► *Colorize* (Figure 8-15). With the Hue, Saturation, and Lightness sliders, you get much more control over the image (with a full-sized preview). The only tricky part is finding the right hue. For sepia, you'll probably want to choose a hue down in the thirties, and then increase lightness as needed to get the right effect.

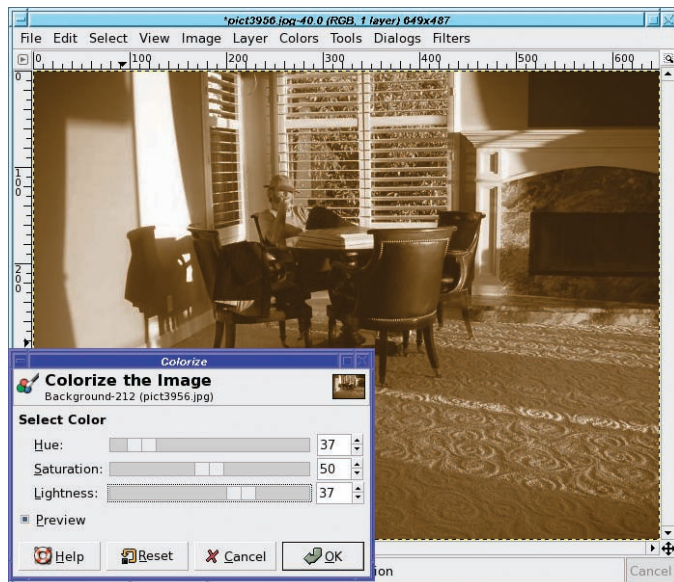



Figure 8-15. Tinting a black-and-white photo with Colorize

## Layer Modes

Finally, you can use layer modes. First, set the foreground color to a medium sepia color: a28a65 (typed into the *HTML Notation* box) works well as a starting point. Now, create a new layer that is entirely filled with this color by clicking on *New Layer* in the Layers dialog. Set the *Layer Fill Type* to *Foreground color*. This covers your original layer. But don't panic!

Go to the *Mode* menu in the Layers dialog and choose *Color* (Figure 8-16).



**Figure 8-16.** With the *sepia* layer selected, choose *Color* in the *Mode* menu.

This gives a sharper result than *Colorize* or *Old Photo*. Of course, you can use the *Color* layer mode to tint photos any shade, not just sepia; and you can create very impressive effects by drawing different colors on different parts of the color layer. You'll learn more about coloring with layer modes in Chapter 10.

## Using Threshold to Clean Up Scans

Have you ever tried to scan in a page from a book? It's especially tricky with an old yellowed book, or a page that has crumples and folds and dirt spots on it, or text showing through from the other side of the page. How can you clean up a page of text like Figure 8-17 into a nice black-and-white image?



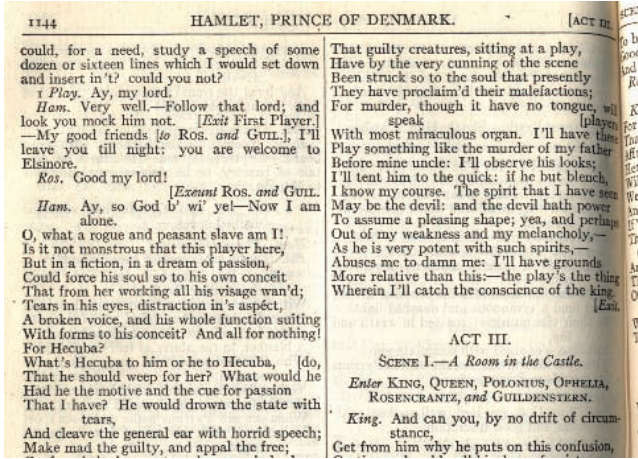


Figure 8-17. A scanned page from an old yellowed book

For anything that really ought to be black and white (with no shades of gray in between), the tool of choice is **Colors** > **Threshold...** (Figure 8-18). Previous GIMP versions have it in **Layer** > **Colors** > **Threshold...**, and it can also be found in **Tools** > **Color Tools** > **Threshold...**

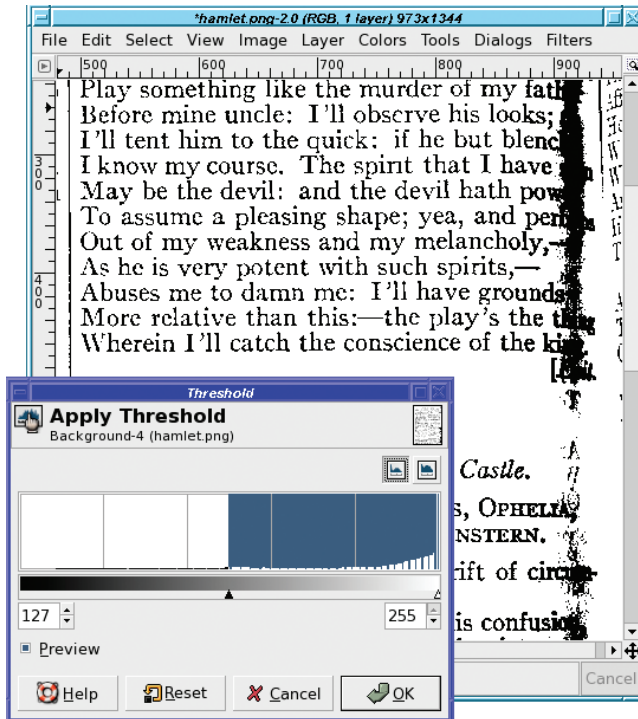


Figure 8-18. The Threshold tool



Threshold has two sliders, though most of the time, you'll only need to adjust the left one. The sliders move along a bar showing a range of grayscales. The key to the Threshold tool: any pixel value between the two sliders—marked as blue in the tool—will be white in the resulting image. Anything that's darker than the left slider or brighter than the right slider will be black. (Yes, that's slightly counterintuitive. To give the image more white, you want less white—more blue—in the tool dialog.)

To adjust the darkness of an image such as a scan, slide the left slider until you're happy with the result.

So what's the right slider for, anyway? Adjusting the right slider will take anything that's very bright in the image, and make it black instead of white. This may sound like an odd thing to do, and generally it's not useful in scans. However, you can also use the Threshold tool on an image to make selection masks. Sometimes being able to select a brightness region somewhere in the middle can help a lot in making a selection. You'll see how later in this chapter.

Threshold, since it creates an image that's purely black and white, will eliminate any "aliasing" smoothing the edges of lines, sometimes resulting in a "jaggy" look. When you want smoother edges and don't mind some gray in the image, Curves or Levels may sometimes be a better tool than Threshold. Try all of them to become familiar with the differences.

## Indexed Color

An image with *indexed* color contains only a fixed set of colors, usually 256 or fewer. You've already met indexed formats such as GIF and indexed PNG in Chapter 2. There, you saw an example of how the number of colors can change the quality of the image.

The list of colors in the image is called the image's *palette*, like an artist's palette holding the pigments to be used in a painting. The file size of an indexed image depends very much on the number of colors in the palette. Indexed formats can be very efficient if they have only a small number of colors. They're ideal for small web images such as icons for buttons or forward/backward arrows.

Every image loaded into GIMP has a color mode, represented in the title bar of the image window. For instance, in Figure 8-19, the title bar shows that the image is in RGB mode.



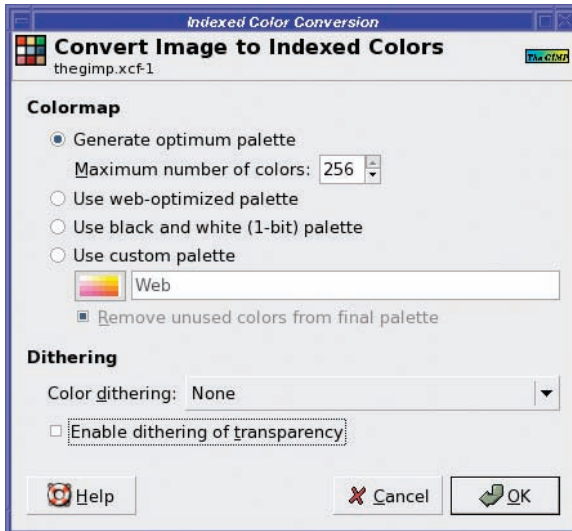
Figure 8-19. A colorful logo

Saving an image as GIF will automatically convert it to indexed mode, since GIF images are always indexed. But you usually won't get the smallest possible image that way: GIMP will use 256 colors, when four might have sufficed. You're much better off converting the image to indexed mode yourself, even for GIF. To save in indexed PNG format, converting to indexed mode is your only option. Since PNG images can be either indexed or full color, GIMP uses the image's mode to determine which form of PNG to use.

## A Typical Indexing Problem: Choosing the Palette

Suppose you've made a logo such as Figure 8-19, and want to save it in an indexed format. How do you make the image size as small as possible?

To convert to indexed mode, use *Image* ► *Mode* ► *Indexed...* Figure 8-20 shows the options.



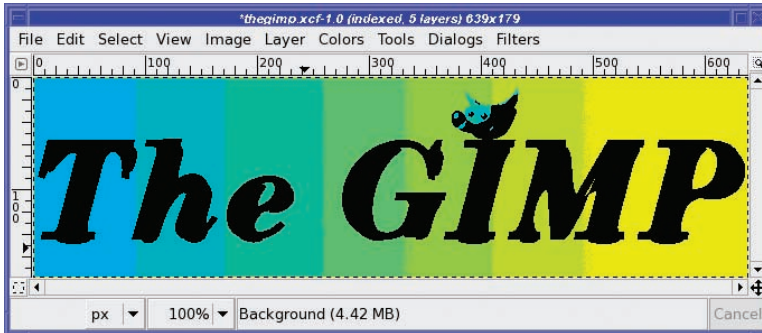
**Figure 8-20.** Options when converting to indexed mode

By default, GIMP chooses *Generate optimum palette*, which creates a palette of the colors in the image, and no others. Most of the time, this is your best choice. The other options offer palettes for special cases.

The black-and-white palette, which contains only those two colors, is useful if you're saving the results of a scan after you've run *Threshold* on it.

The web-optimized palette and the Web palette in the custom palette list use colors that used to be standard in web browsers. Today, nearly all web browsers and computers support 24-bit true color, so web palettes are much less important. You're usually better off optimizing a palette for each image.

The trick in converting to indexed mode is in choosing the right number of colors: the smallest number for which the image still looks good enough. Start with a small number, like four or eight (make a guess based on how many colors you think the image should need), and click OK (Figure 8-21).



**Figure 8-21.** An eight-color palette

*Yuck!* The image looks awful! Time to undo and try a different setting. (Convert to Indexed, alas, does not have a preview.)


## Dithering an Indexed Image

Take a look at those options again. An image that has shades of color can often benefit from *dithering*: combining several colors in the palette to simulate colors that aren't there. Dithering will increase file size significantly, but not as much as increasing the number of colors.

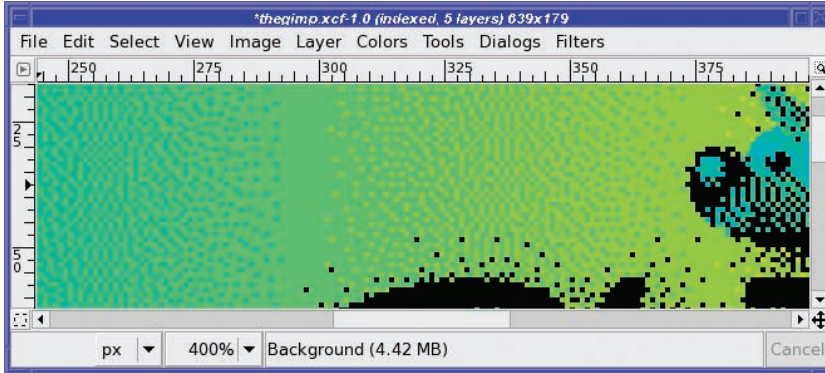
The dialog offers several dithering options:

- The two *Floyd-Steinberg* options differ only subtly.
- *Positioned* dithering is usually best for animated GIFs.
- *Enable dithering of transparency* can help if you have any partially transparent layers.

Figure 8-22 shows the effect of starting with that same 8-bit palette and adding *Floyd-Steinberg (reduced color bleeding)* dithering plus dithering of transparency. The shadows around the letters are now dithered with black dots, but they still don't look at all like a shadow. There's good news too: the transition in the color background looks much smoother now, as you can see in the zoomed image (Figure 8-23).



**Figure 8-22.** An 8-color palette with *Floyd-Steinberg (reduced color bleeding)* dithering and *dithering of transparency*



**Figure 8-23.** Figure 8-22 is zoomed in to show how dithering helps the background look smooth.

### Flatten for Better Dithering

What's the problem? Well, part of the problem is that converting an image with multiple layers to indexed mode often doesn't work very well. This image has five layers: the background, the text, the text's shadow, Wilber, and Wilber's shadow. The GIMP is trying to index each of these layers separately—and so the dithering on those shadow layers looks very poor. There are two solutions.

First, get rid of the layers. *Image* ► *Flatten Image* combines all layers into one. That means it will be difficult to do any further editing on the image, so it may be best to make a copy of the image (*Image* ► *Duplicate*). Then flatten and convert the copy so that you can keep all your layers in your original XCF.

The flattened image looks much better when indexed, even with only eight colors (Figure 8-24).



**Figure 8-24.** Flattening the image helps quite a bit.

### Redesigning for Better Indexed Results

But sometimes, if the goal is to make a very small image, simplifying the design can help more. The sad fact is that those lovely, subtle color fades you can make in GIMP don't translate very well to indexed formats.

If you can change the design to eliminate shading, you can make a *much* smaller image. In Figure 8-25, all the shading (except in Wilber himself) has been removed and the image converted to indexed with no dithering. This resulted in a GIF file less than 5K, as compared to 6.1K for the dithered version and 15.8K for the image from Figure 8-24.



**Figure 8-25.** Shading removed, converted as eight colors with no dithering

Doubtless you're wondering, "Why all this fuss over a mere 10 kilobytes?" But some web pages have hundreds of images on them. If you design your icons to have small file sizes, that can make the difference between a web page that loads instantly and one that loads so slowly that users give up on it and go someplace else.

---

**Tip** Do you ever find that the GIMP function you want to use is grayed out when you're working with an image you loaded from a GIF file? Most often, this is because your image is in indexed mode, perhaps because it was originally a GIF image. Many of GIMP's color functions will only work on an RGB image (or sometimes a grayscale image), and will be disabled for indexed images. (This is because an indexed image does not allow a full range of colors.) If you see color functions disabled and wonder why, check the title bar of the image window and see if it says "Indexed." If so, a quick conversion (using the *Mode* menu) back to RGB mode will solve the problem. (If that doesn't enable the function, check whether the image has an alpha channel; some filters require one.) Of course, if you ultimately want to save to GIF or indexed PNG, you'll have to convert back to indexed mode before you save.

---

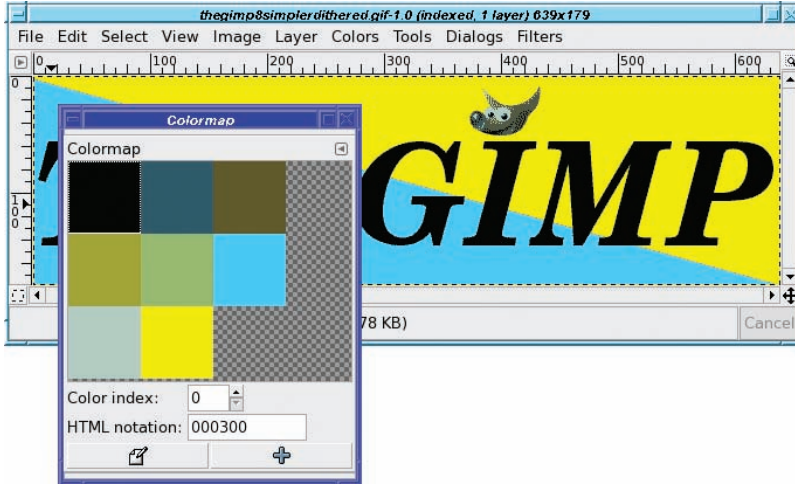
## Editing the Palette

What if you want to change one of the colors in an indexed image?

*Dialogs* ► *Colormap* brings up the colormap editor (Figure 8-26). This window shows every color currently in the image.

To change a color, simply double-click on the color square to bring up a color chooser. Alternatively, click on a square and then edit the *HTML notation* box and hit Return. Or click on a color, and then click the lower-left button in the dialog, *Edit color*.





**Figure 8-26.** *Editing a colormap*

*Color index* lets you move to the next or previous colors, though clicking on the appropriate square is usually just as easy.

The + button in the lower right is *Add color from FG*. It adds a new entry in the colormap, using GIMP's current foreground color. Note that this will slightly increase the size of your colormap, and therefore your image.

Between the colormap editor and making corrections with the Pencil tool, you can usually fix any dithering problems in your indexed image. In the end, you get a clean image to put on the web.

## Pick Colors from the Image

While cleaning up that indexed image, the Pick Colors from the Image or Color Picker tool was awfully helpful. It was introduced back in Chapter 3, but it comes in especially handy here.

One problem you may have with indexed images is difficulty in using GIMP's normal color chooser. For instance, in the indexed Wilber image, if I set GIMP's foreground color to bright red and try to draw on the image, what actually appears is a muddy brown. What gives?

An indexed image can display *only* the colors in the palette. If you try to draw with a color that's not in the palette, GIMP will try to pick the closest color—almost certainly not what you want.

The solution is to use GIMP's Color Picker tool to grab a specific color from the image, as in Figure 8-27. That way, you'll be sure of getting a color from the right palette.



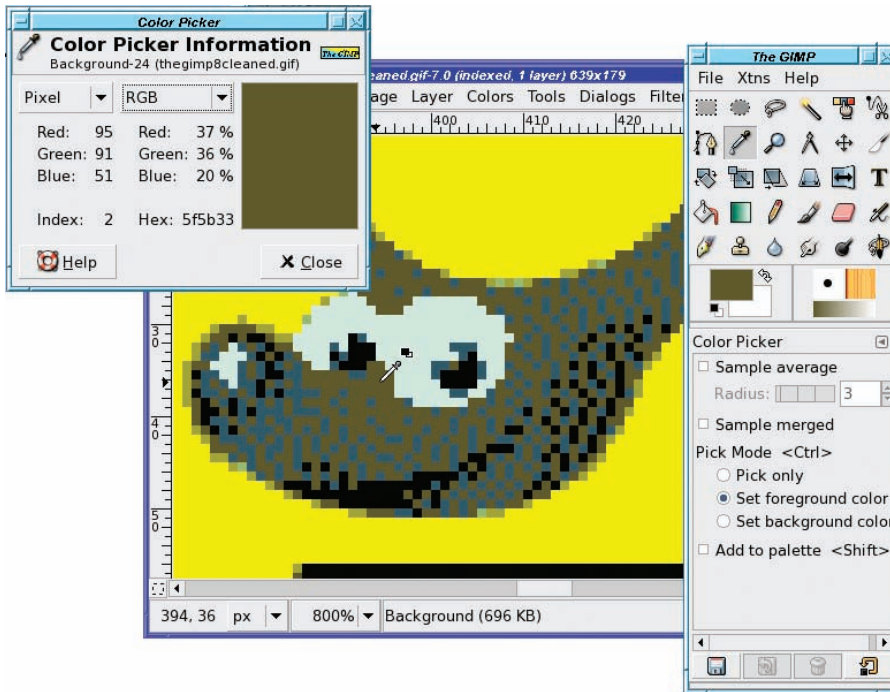


Figure 8-27. Grabbing a color

## The Color Channels

The GIMP maintains three special color masks that show the red, green, and blue components of the current image. You can see them in the Channels dialog, usually docked as a tab with Layers and Paths (Figure 8-28). Light means a lot of the indicated color; dark means little. Notice the white flower in the blue channel, the general lightness of the green channel (there's a lot of green in the photo), and darkness of the red (there's not much red anywhere).

However, the previews in the dialog are too small to see much detail. (You can make them a bit larger using the *Preview Size* item in the tab configuration menu—to the right of the “Layers” label in the dialog—but they’ll still be small compared to the full-sized image.) Also, not many operations work directly on the color channels. You can turn a color channel into a selection (via the context menu in the Channels dialog), but for more detailed work with color channels, you’re usually better off using *Decompose* to create an image with red, green, and blue layers.

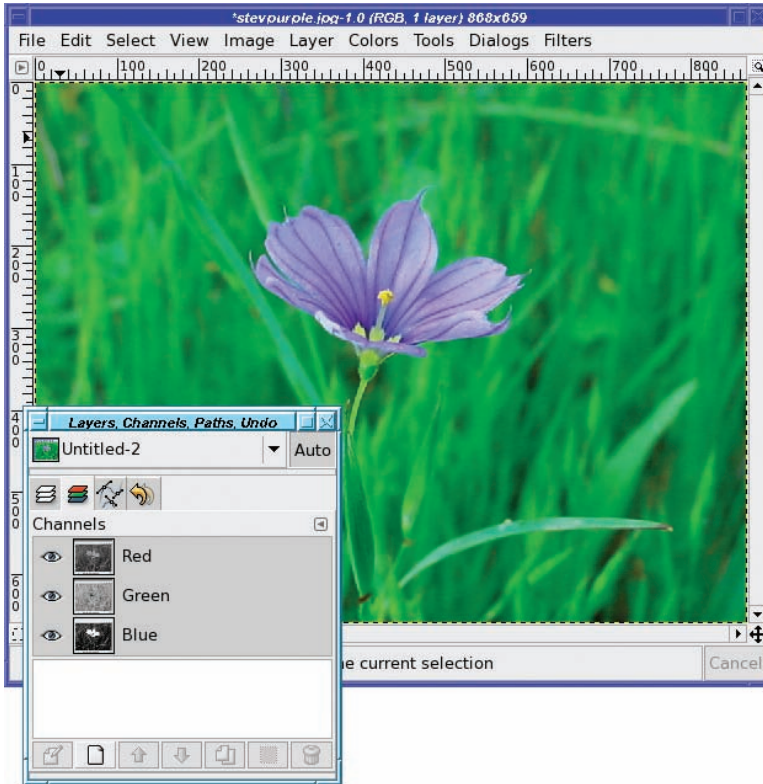


Figure 8-28. The color channels

## How to Use Color Decomposition

Remember Decompose, discussed earlier under “Working with Grayscale or Black and White”? There, you saw how to use it to convert images to grayscale in several different ways. But it’s far more useful than that. What interesting things can you do with color decomposition?


Well, for example, you could select that flower in Figure 8-28.

If you right-click on the blue channel (in the Channel Dialog) and choose *Channel to Selection*, the “marching ants” make it look like you have a pretty good selection of the flower. But that’s deceptive.

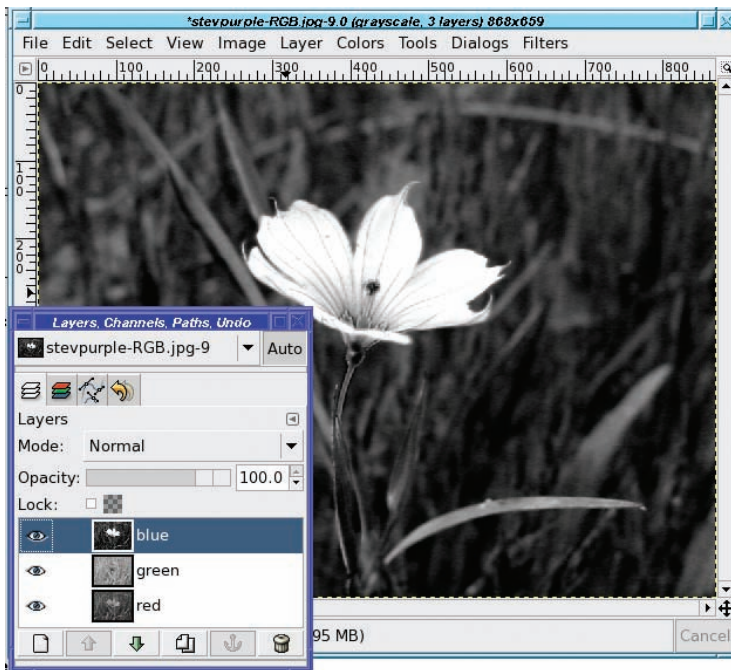
There’s a problem: nothing is *completely* selected. The flower was mostly blue, but not pure blue, so it gets mostly selected. The grass around the flower has some blue, so a little bit of that gets into the selection too. Instead of getting all of the flower and none of the grass, you get a mostly opaque flower and mostly transparent grass.

If you paste the selection onto a white background, you’ll see something like Figure 8-29. It might make a nice background for stationery, but it’s not what you’d want most of the time.

Decomposing to RGB can give a better result. Figure 8-30 shows the blue layer, which is basically the same as the blue channel. You can see lots of pieces of the background, not just the flower. But since it’s in a layer now, you can modify it to be more useful.



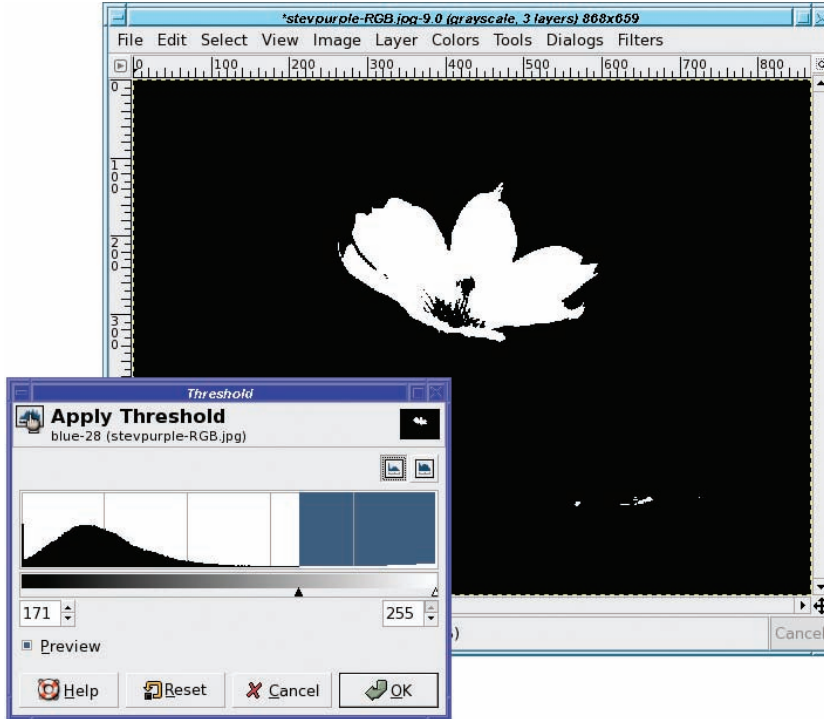
**Figure 8-29.** The flower was only partially selected.



**Figure 8-30.** The blue layer of the RGB decomposition

## Use Threshold and QuickMask to Improve the Selection

To turn the blue layer into a useful selection mask, use the Threshold tool (Figure 8-31). (A tool such as Levels or Curves will also work.)



**Figure 8-31.** *Threshold on the blue channel gives a much cleaner selection.*

Once you've set the threshold, you can use drawing tools to clean up the selection. Paint out (in black) anything you don't want, and paint in (in white) anything you do. Once you're happy with your selection mask layer, copy it (Control+C).

Now go back to your original image, and activate the QuickMask (remember this from Chapter 5?).

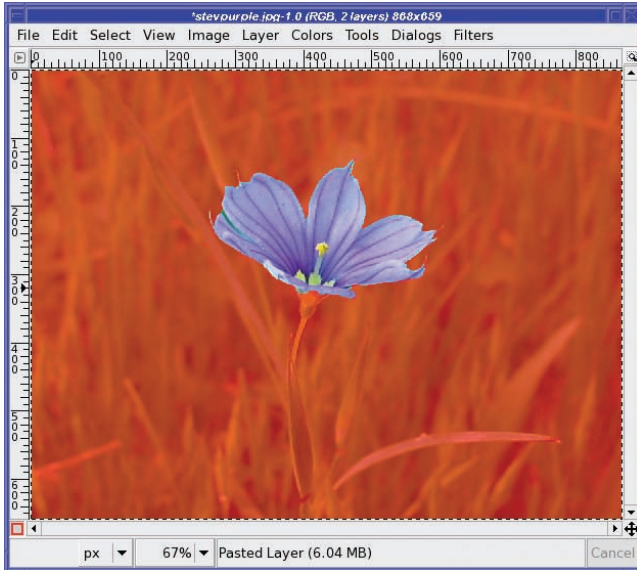
A simple paste (Control+V) onto the QuickMask, and your selection is ready (Figure 8-32). Switch from QuickMask mode back into normal mode, and you have a perfect flower that's ready to paste anywhere.

---

**Tip** This method also works well for creating a layer mask. To avoid the sharp edges that the Threshold tool can create, try Curves or Levels instead.

---

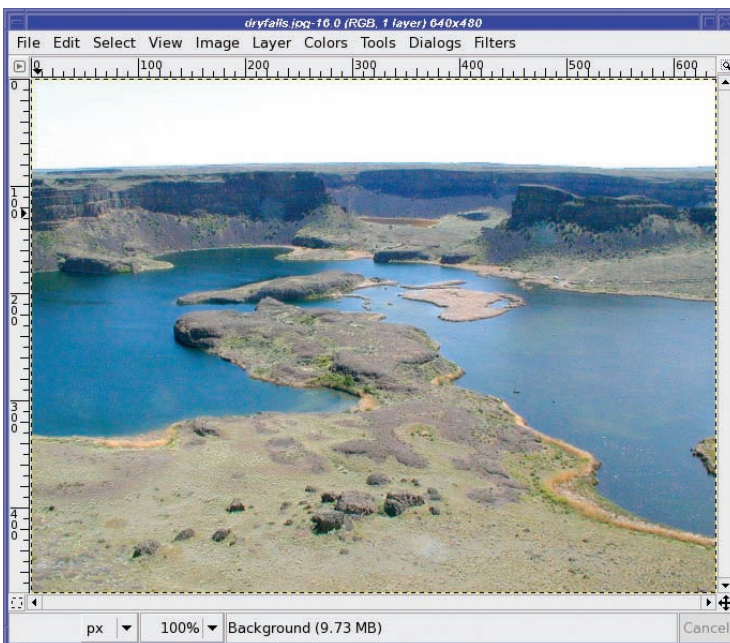




**Figure 8-32.** *The flower is now selected.*

## Decomposing to HSV

With Decompose, though, you aren't restricted to using just color channels. An HSV decomposition can be really useful in extracting a glaring overcast sky (Figure 8-33).



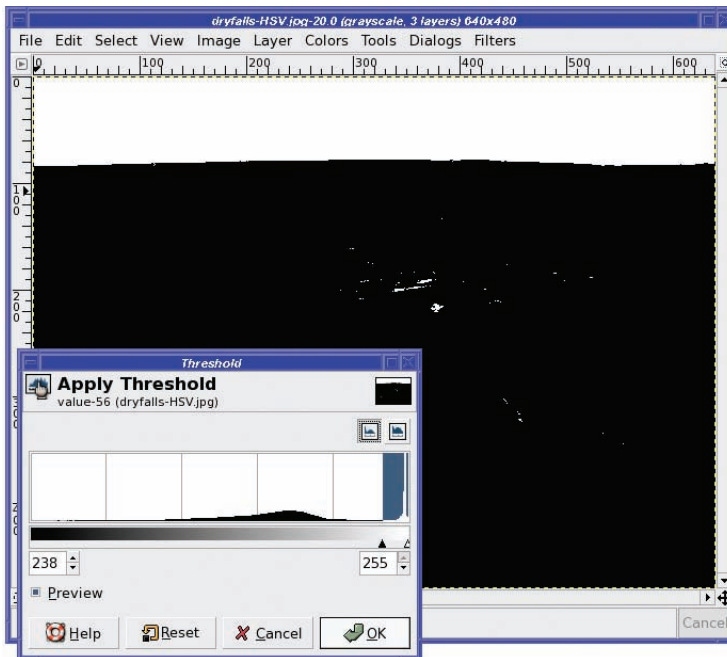
**Figure 8-33.** *The overcast sky ended up overexposed, spoiling this photo.*

Decompose to HSV, and then run Threshold on the value layer (Figure 8-34). Clean up any stray pixels with a drawing tool such as the Pencil, and you're ready to select the sky.

Once the sky is selected (you may also want to *grow* or *feather* the selection by a few pixels), you can make a replacement light blue sky (Figure 8-35). But it doesn't look right.


What's wrong with the sky in Figure 8-35? Real skies are darker and bluer up high, fading to lighter colors down low. So use a gradient and two similar shades of blue to make a more realistic sky (Figure 8-36).

Selecting based on the Saturation channel is also useful. For instance, in a photo showing a blue sky with clouds, the blue sky is more saturated than the clouds, so you can select just the blue sky, and then make it darker blue . . . or green, or red.




**Figure 8-34.** Run Threshold on the value layer.





**Figure 8-35.** Solid-colored skies don't look realistic.



**Figure 8-36.** A gradient between two shades of blue makes a much more realistic-looking sky.

## Color Profiles

Different computer systems display colors very . . . differently. If you've ever viewed the same photo on both a Mac and a PC, you've probably seen how much the colors and brightness can vary.

Some high-end monitors are designed to display colors very accurately. They have a *color profile* describing the relative brightness of red, green, and blue.

High-end printers (like the kind used at professional printing shops) also have color profiles. Since they can produce prints that look very different from the monitor view, professional artists can tear their hair out trying to match the colors they see on a monitor to the colors that will be printed.

If you specify the two color profiles, monitor versus printer, the GIMP can help you predict what a print should look like.

GIMP 2.4 introduced a new category in the Preferences dialog: *Color Management* (Figure 8-37). These preferences give you a lot of control over how colors are displayed. Most people won't ever need to set any of these options, but those who need accurate color can use them to set hardware profiles. If you don't know the color profile for your monitor and printer, you're probably best off leaving the default settings unchanged.

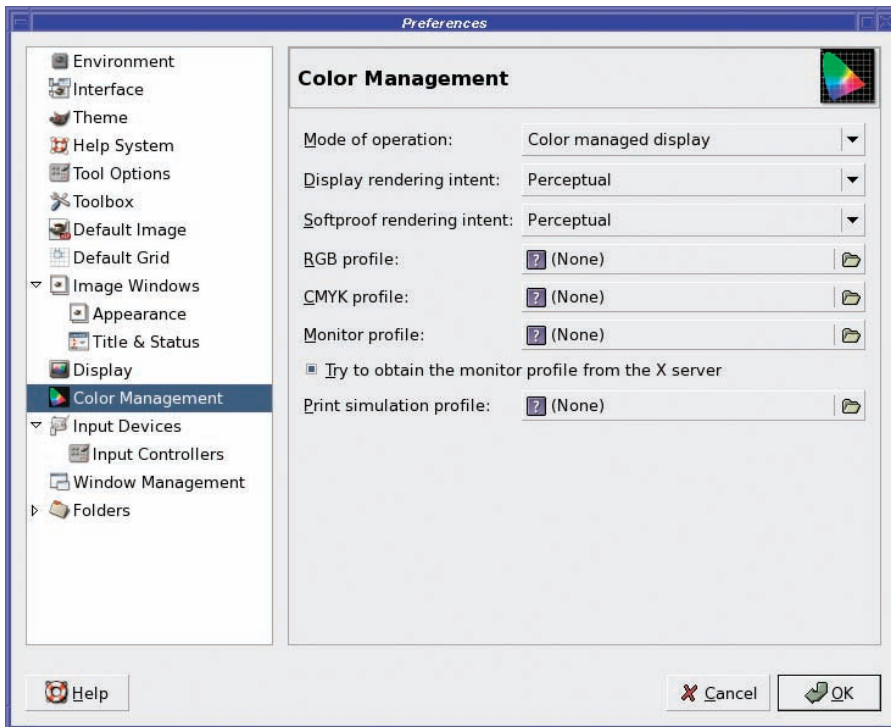


Figure 8-37. The Color Management preferences panel

## Summary

By now, you probably know more about color than you ever wanted to know.

You're familiar with all of the common color models and some of the theory behind them. You know about indexed color, grayscale, RGB, and alpha; how to tell the color mode of your image, and how to change it.

You've worked with the GIMP's different color choosers, and you've probably picked a favorite that you use most of the time. You may even switch off among several of them depending on the sort of project you're working on.

You know several different methods to change a color image to black and white (or simply reduce the amount of color in an image), how they differ, and when to choose each one. You've also seen several ways to add color to monochrome images.

When you need to use an indexed image, you know how to optimize it so you can get a good-looking image that only requires a very small file size. You can clean up scanned images by adjusting their color levels. Finally, you know how to decompose an image into its component colors, and how you can use that to help you in tasks such as selecting a specific part of an image.

You'll find many of these techniques helpful as you move on to the next chapter and learn some advanced drawing tricks.





# Advanced Drawing

In Chapter 4, you learned the basics of drawing with the GIMP. But there are still quite a few tricks that can give your drawings more polish.

In this chapter, you'll learn how to use layer masks, gradients, and layer modes to create great-looking drawings and logos. You'll learn how to control the Perspective tool, and create shadows and highlights. Finally, you'll learn how to create your own brushes, gradients, and other GIMP resources. The chapter will cover:

- Useful mask tricks
- Layer modes
- Creating depth: Drawing with layer modes
- Drawing realistic shadows
- Realism and multipoint perspective
- Adding reflections and shading
- Creating brushes, patterns, and gradients

## Useful Mask Tricks

You learned about the basics of layer masks in Chapter 5, and saw one use for them. But there are all sorts of interesting effects you can create with layer masks.

### Making Text “Fade Out”


One of the easiest uses of a layer mask is to make some text (or any other layer) fade out.

Just follow two easy steps: add a layer mask to the text layer, and then draw a gradient on the text layer:

1. Add a layer mask. Right-click on the layer in the Layers dialog and choose *Add Layer Mask...* Initialize it to *White (full opacity)*, the default setting.
2. Draw a gradient. The layer mask should already be selected since you just created it. If not, click on its thumbnail in the Layers dialog. Choose the Gradient tool in the Toolbox, and the standard foreground and background colors of black and white, respectively.

Then just drag in the image to draw the gradient. You can drag exactly vertically or horizontally (remember, holding down the Control key while you drag will help by constraining the gradient's angle to a multiple of 15 degrees), or you can drag in any direction you like.

In this case, I started with a layer that had text and drop shadow already merged together, and then dragged diagonally, from the lower-right toward the upper-left of the text-and-shadow layer (Figure 9-1). You can see the gradient in the layer mask's thumbnail in the Layers dialog.



**Figure 9-1.** A gradient on the layer mask makes the text fade.

## Making a Fuzzy Border

A layer mask can also give an image a fuzzy border, so that it fades into the background. This gives a nice effect for images included on web pages.

The procedure is very straightforward. First choose your favorite selection tool. Turn on feathering (you'll probably want as much feathering as possible), and then make a selection where you want your border to be (Figure 9-2).





**Figure 9-2.** Make a selection where you want the border to be.

Now add a layer mask to the image layer: right-click on the layer in the Layers dialog and choose *Add Layer Mask...* (In some GIMP versions, you may not be able to create a layer mask unless the image has transparency. If *Add Layer Mask...* is grayed out, use *Layer > Transparency > Add Alpha Channel* for that.) Under *Initialize Layer Mask to:* choose *Selection* (Figure 9-3). This will make the new layer mask white inside the selection and black outside of it, which is exactly what you want.

Now you have an image with a fuzzy boundary. Everything outside the selection is transparent, so GIMP will show it with a checkerboard pattern.

To find out how the image will really look on a web site, there are two additional steps. First, hide the selection with *View > Show Selection* or *Control+T*. Second, make a new layer of white (or any other background color you prefer) and move it below the image layer. That will give you a much better idea of how your fuzzy border looks (Figure 9-4).

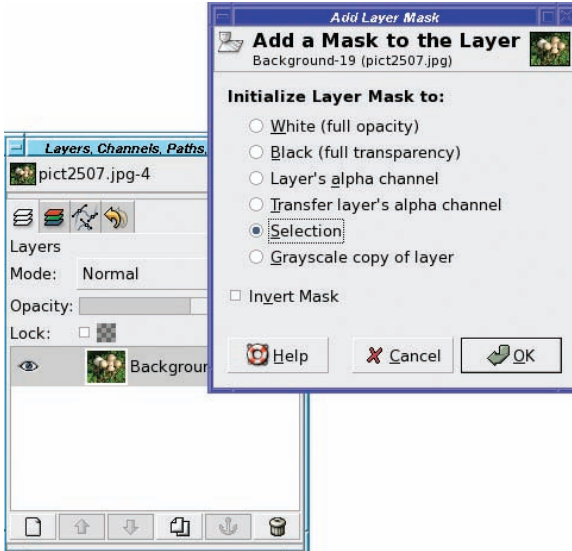


Figure 9-3. Make a new layer mask, initialized to “Selection.”

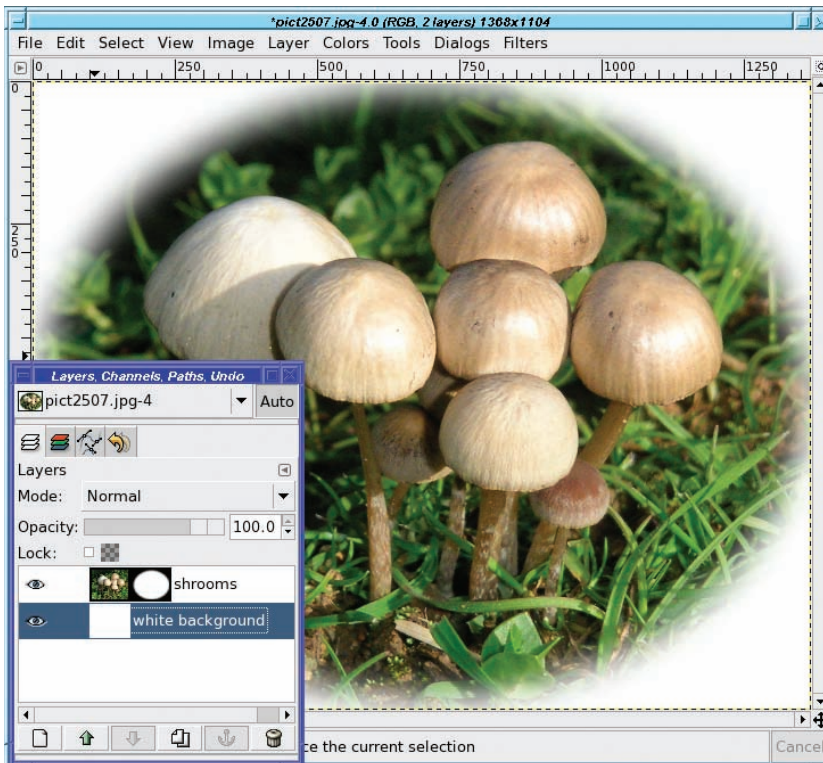


Figure 9-4. Adding a white background layer and making the selection invisible gives you a much better idea of how the border will look.


If you're happy with the result, make the selection visible again (that way you won't be surprised the next time you try to select something), and then cancel it (*Select* ► *None*). Save the image as JPEG since it's full color—unless you absolutely need the transparency. If you do, make the background layer invisible (you won't want to save that part), and then save in a format that allows transparency, such as GIF or PNG (refer back to “Image File Types” in Chapter 2 for a discussion of the trade-offs).

## Even Bigger Fuzzy Borders

What if you want more feathering than the 100 pixels that the GIMP allows in its selection tools?


Well, while you still have the selection, you can use *Select* ► *Feather...* But that doesn't let you preview what the border will look like.

A better solution is to make sure the layer mask is selected (it will have a white border, as in Figure 9-5)—and then blur it with *Filters* ► *Blur* ► *Gaussian Blur...*



**Figure 9-5.** A selected layer mask has a white border.

Blurring the existing 100-pixel feathering by another 250 pixels gives a more gradual border, as shown in Figure 9-6.



**Figure 9-6.** Blur the layer mask to get a more gradual border.

But wait—what's wrong with Figure 9-6? Notice what happens at the edges. All this blurring has made the border extend beyond the edges of the image. If you put the image on a web page, you'll see squared-off edges, as in Figure 9-7.





**Figure 9-7.** *Blurring beyond the size of the image creates squared-off edges.*


This happens because the oval frame gets bigger when it's blurred, and the image isn't big enough to contain it. For a really fuzzy border, you have to leave quite a bit of room around your original selection. If you think you might want a very fuzzy boundary, you're best off not cropping much until the very end.

---

**Note** You can do the same sort of operation on GIMP's QuickMask: select the image with a fuzzy border, and then paste it onto a white background. Why use a layer mask? Because that makes it easier to see the effect of the operation. With the QuickMask, you have to modify the mask, change it to a selection, copy it, and paste it somewhere else—only then will you see whether you need to go back and modify the QuickMask.

---

Of course, since the layer mask is just a black-and-white image itself, you can modify it in a variety of ways to make all sorts of borders, like the ones in Figure 9-8. Experiment! You may not like most of the effects, but you might also find the perfect look.



**Figure 9-8.** You can make all sorts of frames using layer masks.

## Interpreting the Mask Border Color

Sometimes it's hard to tell the mode of a layer mask. But it's a lot easier if you learn to decode the colors of the mask thumbnail's border in the Layers dialog:

- *Black*: Not selected.
- *White*: Selected. If you draw, you will draw on the mask.
- *Green*: Displayed. GIMP will display the contents of the layer mask, rather than applying it as a mask. When in this mode, the border color doesn't help you tell whether the mask is selected for drawing; to find out, check the color of the layer preview next to the mask preview. If the layer preview is black (unselected) but the layer is highlighted in the dialog, the layer mask is selected.
- *Red*: Disabled. Again, you can't tell whether or not the layer mask is selected if it's disabled, so check the layer preview's border color.



## Layer Modes

By now, you should be very comfortable with adding layers to an image. Generally, when you put one layer on top of another, the top layer is what you see. Lower layers only show through when the upper layer is at least partially transparent.

But layers can be combined in many other ways, using *layer modes*. You'll find them in the *Mode* menu in the Layers dialog. GIMP will use a layer's mode to determine how to combine each pixel from the top layer with the pixel in the same location in the layer below. Layer modes are also called *blending modes*, because they specify how to blend the colors in a layer with the colors in the layer below it.

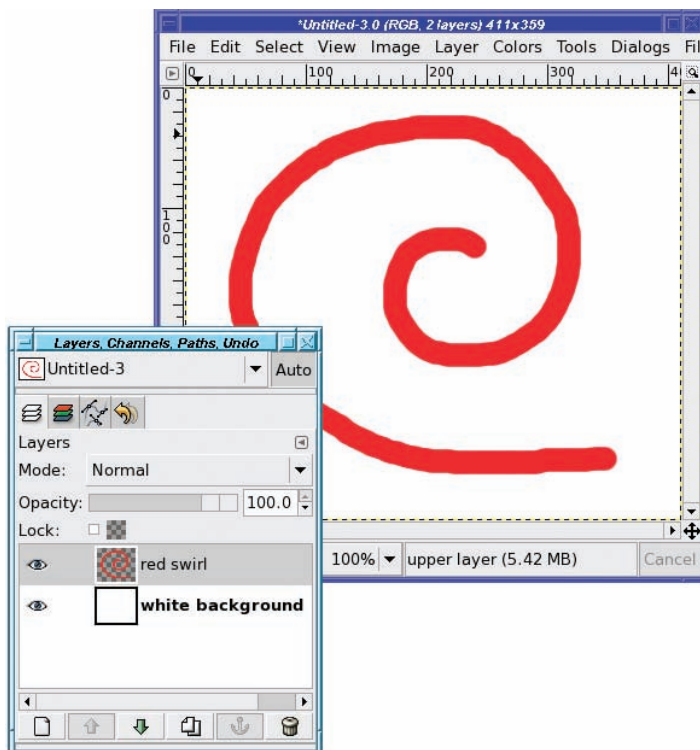
---

**Note** The *Mode* drop-down menu in the drawing tool options, which you've already seen in Chapter 4, works basically the same way that layer modes do. It also adds a few special modes that are only used for drawing tools.

---

Confusing? Let's look at an example.

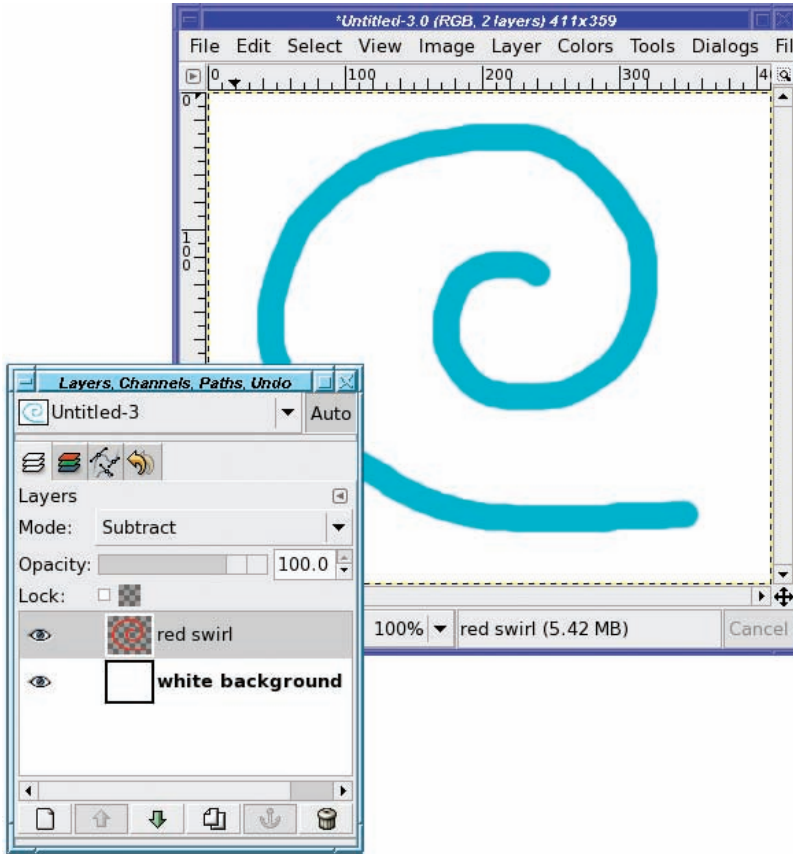
Figure 9-9 shows a red swirl on a white background. Of course, the swirl is in a separate layer from the background.



**Figure 9-9.** A red swirl layer on a white background layer

Where the swirl is red, you're seeing just that—red. The layer mode is *Normal*, so GIMP takes each pixel in the top layer (the red swirl), determines that it's not transparent, and shows you only that. Outside of the swirl, the top layer is transparent, so GIMP shows you the next layer down. Simple!

But what happens if you set the layer mode to *Subtract* (Figure 9-10)?



**Figure 9-10.** The red swirl with *Subtract* mode

Why did the red swirl turn that color?

In *Subtract* mode, the color you see is the result of subtracting the front layer (the swirl) from the back layer (the background). Each color channel is subtracted separately.

Remember that white is the combination of the three primary colors: red, green, and blue. Each color has a range from 0 to 255 (8 bits, in computer terminology). A white pixel is (255, 255, 255) to the GIMP. A red pixel is (255, 0, 0): all the red, with no green or blue. So:

$$\text{white} - \text{red} = (255 - 255, 255 - 0, 255 - 0)$$


or (0, 255, 255).

The result has no red at all, but full green and blue: it's cyan. That's why the red swirl on top of a white background becomes cyan in *Subtract* mode.

A lot of the modes don't do anything interesting when the background is white. For instance, in *Addition* mode, the colors of the two layers are added together. But white already contains a full complement of each color. You can't get any whiter than white, so in Addition mode, the red swirl just becomes white and disappears in the background. For the same reason, other modes won't seem to do anything useful if the background is black.

## A Quick Tour Through All of the GIMP's Layer Modes

Figure 9-11 shows all of the GIMP's layer modes, using the tree from Chapter 4 against a background of light blue/cyan.



**Figure 9-11.** All of the GIMP's layer modes

Be careful not to draw too many conclusions from this one diagram, though. The effects of each layer mode can vary considerably depending on the colors of the layers involved, as you'll see. Also, though it may not seem like it as you read through, all these modes have their uses—which are often not what you'd expect at all! But more on that later.

Dissolve and the two special drawing tool modes Behind and Color Erase were discussed in Chapter 4. Now I'll cover the other modes, what they do, and some ways to use them. You'll learn even more uses for the layer modes in Chapter 10.

### MATHEMATICALLY SPEAKING

In the discussions of the individual modes, I'll give equations for what each mode is actually doing “under the hood.” For the mathematically inclined, this can help to understand better what the modes do.

If that doesn't describe you, don't panic! You don't have to understand the equations to “get” layer modes. Many talented artists use the GIMP's layer modes to create beautiful art without ever worrying about the mathematical underpinnings. If you're not comfortable thinking about colors as vectors, or with the mathematical operations, just skip the equations (and the rest of this sidebar).

When equations are given, rather than using the (R, G, B) notation used above, I'll use a shorthand notation, where each pixel is referred to as a single vector with three color components. The two layers will be referred to as F, for the layer in *front*, and B, for the layer in *back*. For instance, in the Subtract mode example above, a shorter way of writing the operation is:

$$B - F$$

This is just a shorthand way of saying “For each of the three colors in B and F, subtract the F color from the corresponding B color.”

Colors must always range from 0 to 255. If not otherwise stated in an equation, you can safely assume that if the result of an operation results in a negative color, 0 will be substituted; if an operation results in a value greater than 255, GIMP will use 255. So the result of subtracting bright red from medium gray is:

$$(128, 128, 128) - (255, 0, 0) = (0, 128, 128)$$

or, no red, half green, and half blue: a dark cyan.

## Addition, Subtract, and Difference

*Addition* mode is straightforward: the two pixel values are added together.

$$F + B \qquad \qquad \qquad (Addition)$$

Of course, if that makes any of the colors greater than 255, then that color is set to 255. Addition will nearly always make the image lighter.

*Subtract* is the opposite: the front layer is subtracted from the layer below it.

$$B - F \qquad \qquad \qquad (Subtract)$$

This will make colors darker, naturally. You'll often end up with a lot of black or near-black from using Subtract mode.

---

**Note** Why is it “Addition,” but “Subtract” rather than “Subtraction”? Historical reasons. Which usually means, “No one remembers any more why they're named that way.”

---

Taking the *difference* of two numbers is the same as subtracting, isn't it? Usually, that's true. But in the GIMP, *Difference* mode is, well, different. *Difference* is a subtraction too: but after subtracting, if any color values are negative, GIMP uses the absolute value of that color instead of truncating it to zero.

$$| B - F | \quad (\textit{Difference})$$

This means that instead of a lot of dark areas, the regions that would have been black become lighter again—especially in the colors that would have been absent. You can see that in Figure 9-11: areas of the trunk that were dark blue in *Subtract* mode become magenta in *Difference* mode.

*Difference* sometimes creates strange, garish colors and unusual transitions between colors. It can be useful when you need that sort of effect. In Chapter 10, you'll see several other good uses for *Difference* mode.

## Multiply and Divide

*Multiply* multiplies the value of the pixels in the two layers (bet you already guessed that). This generally makes the image darker.

Why should that be? Multiplying two numbers should make a bigger number, shouldn't it? And bigger numbers are lighter?

The trick is that the result is *normalized* to 255 (in other words, divided by 255). Otherwise, too much of the image would end up being too bright.

So, the actual equation for *Multiply* is:

$$F * B / 255 \quad (\textit{Multiply})$$

---

**Note** This ends up being darker than the front layer because  $B / 255$  is less than 1 (except where B is white). So when it's multiplied by F, you get a number that's usually less than F.

---

*Divide* is a little more complicated. As you learned in grade school, you can't divide by 0 (at least without getting infinity). In order to guard against division by 0, GIMP adds 1 to the front layer's pixel values. Like *Multiply*, *Divide* is normalized, though this time it's normalized to 256 (1 is added to the normalization because of the 1 that was added to the front layer's pixels).

$$B * 256 / (F + 1) \quad (\textit{Divide})$$

Even after normalizing the result, *Divide* mode often results in nearly white pixels that look “burned out.”

## Dodge and Burn, Screen and Overlay

These four modes are closely related to *Multiply* and *Divide*. *Burn* and *Overlay* are similar to *Multiply* and make the image darker; *Dodge* and *Screen* are more like *Divide*, and make an image brighter.

Dodge divides the back layer by the inverse of the front layer (subtracting the front layer from white, just like *Colors* ► *Invert* would do). This usually lightens the image about the same amount as Divide mode does, though sometimes it doesn't whiten the image quite as much . . . but it inverts some colors. (Unfortunately, that's hard to tell from Figure 9-11.)

$$B * 256 / (256 - F) \quad (\textit{Dodge})$$

In Burn mode, it's the back layer that is inverted rather than the front layer; then after the division, the result is inverted again.

$$255 - ((255 - B) * 256) / (F + 1) \quad (\textit{Burn})$$

Burn ends up looking a lot like Multiply, though there are some differences. In Figure 9-11, you can see that they both got darker, but Burn took on the blue color of the background layer, while Multiply turned black. With other color combinations, though, Burn might end up darker.

In Screen mode, the values of the visible pixels in the two layers are inverted, multiplied, and the product inverted again.

$$(255 - ((255 - B) * (255 - F))) / 255 \quad (\textit{Screen})$$

This sometimes results in a “washed out” look and a bright image with muted colors.

Overlay is a combination of Multiply and Screen modes. The back layer is inverted, multiplied by twice the front layer, added to the original rear layer, and then multiplied by the original rear layer. (If you didn't follow that, don't worry about it.)

$$B * (B + (2 * F * (255 - B)) / 255) / 255 \quad (\textit{Overlay})$$

It usually doesn't darken nearly as much as straight Multiply, though in this case, it has taken on so much of the background color that you can scarcely see the tree any more.

## Hard and Soft Lights

*Hard light* is another combination of Multiply and Screen modes. Depending on the color of the rear layer, you can sometimes think of it as an opposite of Overlay mode.

The formula is complicated because it has two parts: one for dark colors, and another for bright ones.

$$\begin{aligned} &\text{if } (F > 128) && (\textit{Hard light}) \\ &\quad 255 - ((255 - B) * (255 - (256 * (F - 128)))) / 256 \\ &\text{else} \\ &\quad (B * F * 2) / 256 \end{aligned}$$

Hard light can sometimes be a useful mode for combining two photographs when you want bright colors and sharp edges.

*Soft light* is not the opposite of Hard light, but it does tend to soften edges and dim colors. It's yet another mixture of Multiply and Screen modes, and its effect is most similar to Overlay mode. In fact, in some versions of GIMP, Soft light and Overlay modes are identical.

If *Rm* is the result from Multiply mode, and *Rs* is the result from Screen mode, the equation for Soft light is:

$$((255 - B) * Rm) + (B * Rs) / 255 \quad (\textit{Soft light})$$



## Darken or Lighten Only

These two are easy: they do just what their names imply. For each color component of each pixel, GIMP will show you the darker (or lighter) of the two components. So *Darken only* will make the image darker, while *Lighten only* will make it lighter. So Lighten only would be:

```
if (B > F)                                (Lighten only)
    B
else F
```

To get Darken only, just reverse the greater-than sign to a less-than sign.

## Grain Extract and Grain Merge

*Grain extract* is supposed to produce something akin to the “grain” you might see in photographic film. In reality, I find it’s more useful for giving images an embossed and desaturated look.

Its equation is surprisingly simple:

```
B - F + 128                                (Grain extract)
```

*Grain merge* adds texture to an ordinary image, as if adding film grain or texture from an artist’s canvas. The equation is the mirror of Grain extract:

```
B + F - 128                                (Grain merge)
```

## Hue, Color, Saturation, and Value

These modes are straightforward once you understand what they do, but their names can be a bit confusing.

*Hue* mode takes the hue from the front layer (unless the front layer’s saturation is 0), but takes saturation and value from the back layer. *Saturation* and *Value* modes, likewise, take the named value from the front layer and everything else from the back layer.

---

**Note** Some graphics programs offer a *Luminosity* mode instead of Value. It’s the same thing.

---

*Color* mode gives you the color of the front layer with the lightness of the back layer. It’s like Hue, except that it uses both hue and saturation from the front layer, combined with the value from the back layer.

You can use this mode to colorize an image, using a photograph as the back layer and a solid color as the front layer. It’s as if you took a black-and-white version of the back layer and covered it with a colored filter—only it works better than simply using a translucent-colored layer in normal mode, since that would darken the image unacceptably.

Saturation and Value modes aren’t needed very often (Saturation can be used to remove color from selected parts of an image), but Hue and Color modes can be quite useful for colorizing images, as you’ll see in Chapter 10.

## Creating Depth: Drawing with Layer Modes

So far, layer modes probably seem mostly an academic exercise. But there are some surprisingly interesting things you can do by combining them.

For instance, suppose I start with a happy face drawn with circular selections, stroking, and the Paintbrush tool (Figure 9-12). Of course, in keeping with the first law of drawing (*Make a new layer!*), the face is on a layer by itself, separate from the white background.



**Figure 9-12.** *A simple happy face*

Make a copy of the face layer, and offset it a few pixels up and to the right. All you'll see initially is a face drawn with slightly thicker lines.

But look what happens when you use layer modes on the top copy of the face (Figure 9-13).




**Figure 9-13.** *The effect of different layer modes on a duplicate of a layer, shifted over slightly*

As you can see, layer modes can give some nifty three-dimensional effects to simple lines with very little effort on your part.

But it gets even better! Layer modes control how a layer combines with everything visible below it—not just the next layer down. In the case of the two face layers, each transparent along the borders of the lines, the color of the background makes a difference.

Figure 9-14 shows what some of the modes look like if the background layer is changed to a medium gray: (128, 128, 128). The effects are even more interesting than the ones on a white background—and very different.



**Figure 9-14.** Layer mode depth effects are even more pronounced when using a background other than white.

---

**Tip** Medium gray is often useful when combining layers to get depth effects. Medium colors can work well too. Blurring one of the two layers can make the effects even more pronounced.

---

By now you're probably worried. "How can I possibly remember what all the different layer modes will do with every possible color combination?" Fear not. You don't necessarily need to remember all the possible combinations. Just make two similar layers (one a copy of the other, but shifted, blurred, or otherwise modified), and then start trying layer modes until you see an effect you like.

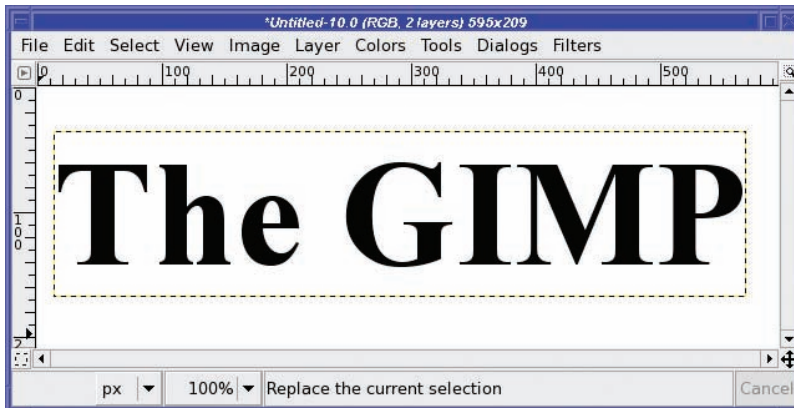
---

**Tip** There's a problem with using a gray background. Ultimately, you'll probably want your pretty 3-D image on a layer of its own, so that you can paste it into other projects, or on a nice color that works for the project. Since the background color is inherent to the effect, you can't just make the background invisible: if you do, the interesting effect will disappear too. (Try it!) The solution: use GIMP's Select by Color tool to select the background color, and nothing else. Then *Edit* ► *Clear* (or invert the selection and then copy your creation). If your object has fuzzy edges, you may need to play with the Threshold, and then clean up a few extra gray pixels that didn't get properly selected . . . but that's a small price to pay for nifty jellybean effects!

---


## Combining Layer Modes: Making 3-D Letters

Here's a straightforward example of how to use layer modes to create interesting logos. Start by making some text in basic black on a white background (Figure 9-15).



**Figure 9-15.** Start with some black text.

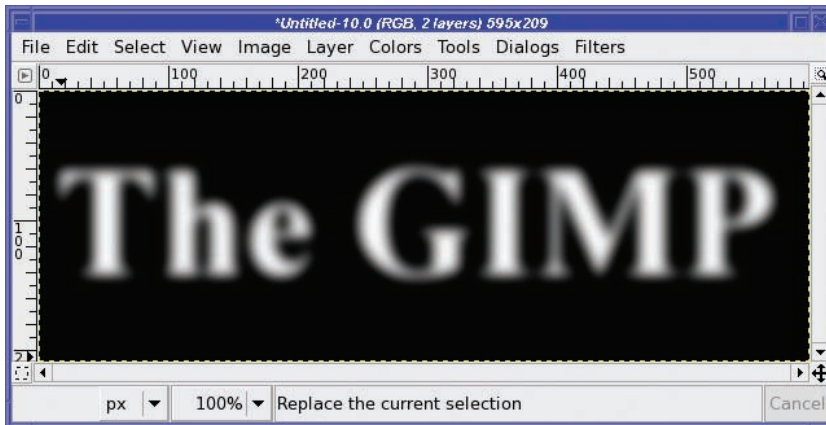
For once, rather than keeping the text layer separate, in this exercise you should merge the two layers. Right-click on the text layer in the Layers dialog and select *Merge Down* (Figure 9-16).



**Figure 9-16.** Merge the text layer into the background.

Make a duplicate of the text layer, Gaussian Blur it (*Filters* ► *Blur* ► *Gaussian Blur...*), and then invert it (*Colors* ► *Invert*) to get Figure 9-17.





**Figure 9-17.** The duplicate text layer has been blurred and inverted.

For the next step, you need to replace the original white background on the bottom layer with a medium gray. (Remember our old friend medium gray?) Turn off visibility temporarily on the blurred and inverted layer so that you can see what you're doing on the lower layer. Then select a medium gray foreground color: in the color chooser, set H and S to 0 and V to 50, or type 808080 into the *HTML Notation* box.

In theory, you should be able to use the GIMP's Select by Color tool to select all the white background, and then use *Edit* ► *Fill with FG Color* to turn it gray. Unfortunately, this isn't enough. When you made the text, you probably used *antialiasing*. (Remember that term from Chapter 3?)

Antialiasing smooths the text by making pixels at the edge slightly transparent, blurring the interface of type and background. That means that along the edges of the letters, the color isn't exactly black or exactly white, but some in-between shade of gray. Select by Color will leave some of these areas unselected, and when you fill with gray, you'll end up with speckles left over (Figure 9-18).

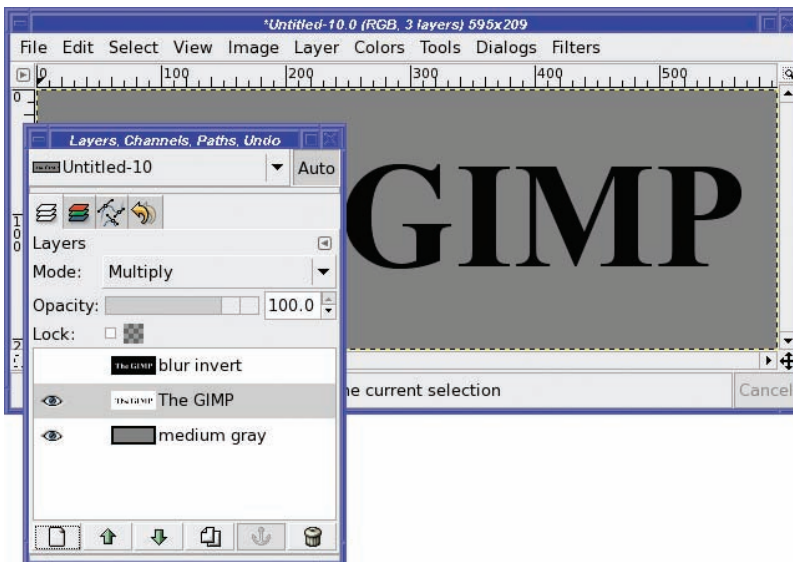


**Figure 9-18.** Light speckles left behind when you try to use *Select by Color* on antialiased text

You could clean those up by hand, or by using operations like *Select* ► *Blur* or *Select* ► *Grow*, but that's a lot of work.

So what *can* you do? Layer modes to the rescue! Make a new layer with the *Layer Fill Type* set to *Foreground color*, so you have a layer of solid gray. Use the down arrow in the Layers dialog to move that layer to the bottom of the stack, below your black-text-on-white-background layer.

Then set the text layer mode to *Multiply*, and magically, the white background turns to gray (Figure 9-19).

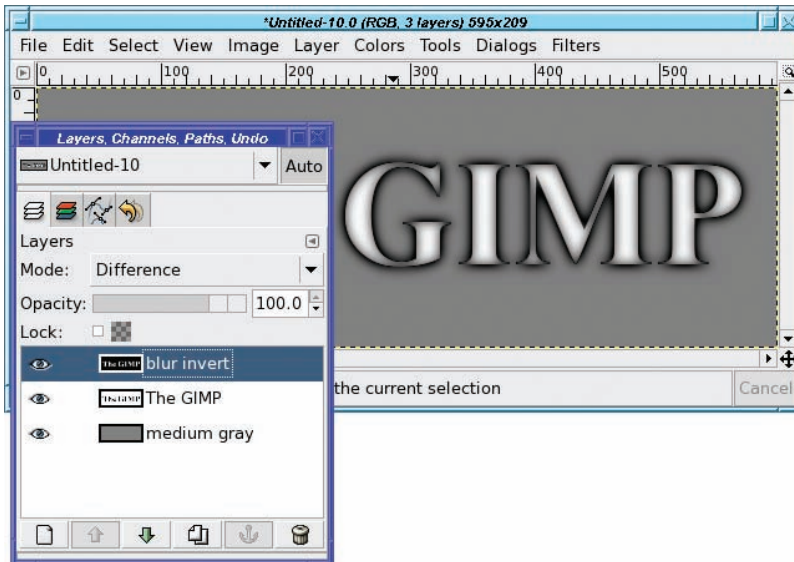


**Figure 9-19.** *Multiply* mode changes the white background to gray smoothly, without any unwanted speckles.

Why does that work so well? How does Multiply get rid of the speckles so cleanly?

Well, it's not actually removing any antialiasing speckles. Instead, it's darkening the white to become the background color, without changing the black text (or its gray antialiasing) very much. Darken only works as well as Multiply, and the trick works for colored backgrounds as well as gray ones. On a gray background like this one, you could also use Burn; but with a colored layer underneath, Burn will change the color of the black text. (As always, if you're not sure, try several different layer modes and see what works best.)

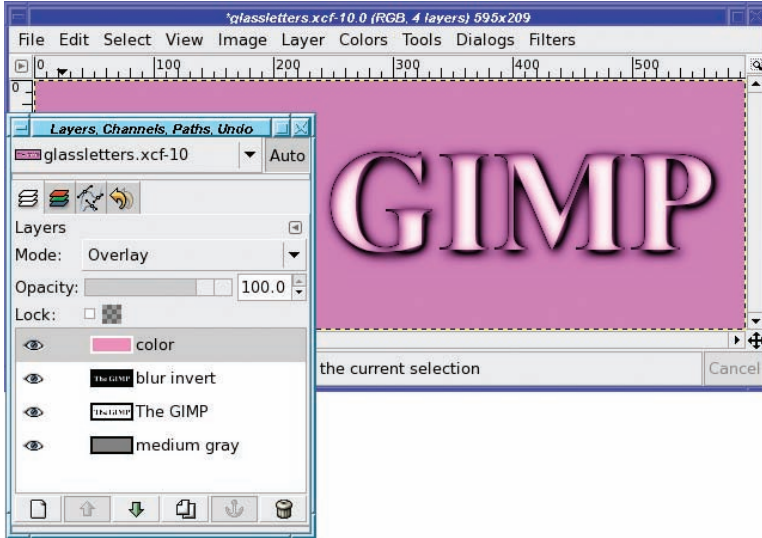
Getting back to the 3-D letter project: turn visibility back on for the top layer (the one you blurred and inverted). Then set that layer to *Difference* (Figure 9-20).



**Figure 9-20.** Set the top (blurred, inverted text) layer to *Difference* mode.

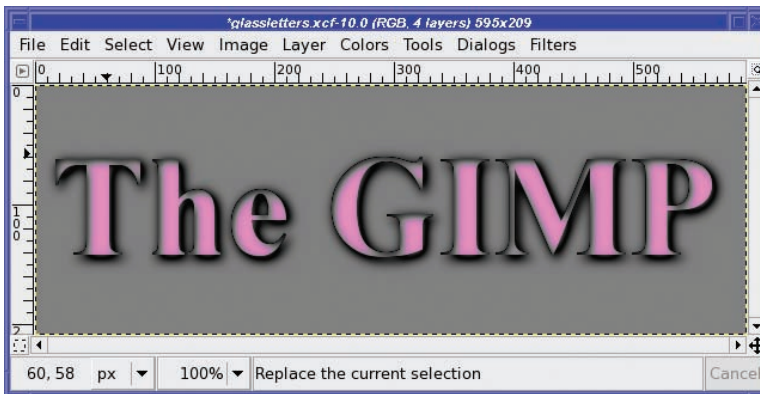
Nice effect! But you can get an even nicer, more three-dimensional effect by using the Move tool to move the top (blurred, inverted) layer a few pixels in any direction.

Don't like the gray color? No problem. Make another solid-color layer, any color you like, put it on top, and set it to *Overlay* mode (Figure 9-21).



**Figure 9-21.** Add a color overlay.

Or set it to Darken only, which maintains a medium-gray background while turning the text your desired color (Figure 9-22). That might be useful if you want to select the text and paste it somewhere else.



**Figure 9-22.** Try Darken only mode for the top color layer.

---

**Tip** Try all the available modes for your color filter. Quite a few of the modes give interesting results.

---

I'm sure by now you realize that I've only given you a tiny idea of the useful effects you can get by combining modes. It would be an epic task to list all the possible combinations. So, make lots of layers and experiment for yourself!

## Drawing Realistic Shadows

Remember adding a shadow to the tree-in-a-box in Chapter 4? Did you find that you had to try over and over to get a shadow that looked good?


There's an easier way to give a complicated object a shadow than by drawing it by hand. Just make a copy of the object's layer, turn it black, blur it, and then use the Perspective tool to make it point in the right direction.

Pasting an object into a photograph is a great way to show this. First, make sure you know in which direction the shadows are pointing in the original photo. In Figure 9-23, that's easy because of the figure of the hiker. I've livened up the picture by adding Tux, the Linux penguin.

---

**Note** Tux was drawn by Larry Ewing using GIMP version 0.54. Larry has an excellent “Penguin Tutorial” page describing the methods he used: <http://www.isc.tamu.edu/~lewing/linux/notes.html>.

---



**Figure 9-23.** Notice the direction of the hiker's shadow. Tux looks somehow artificial without a shadow.

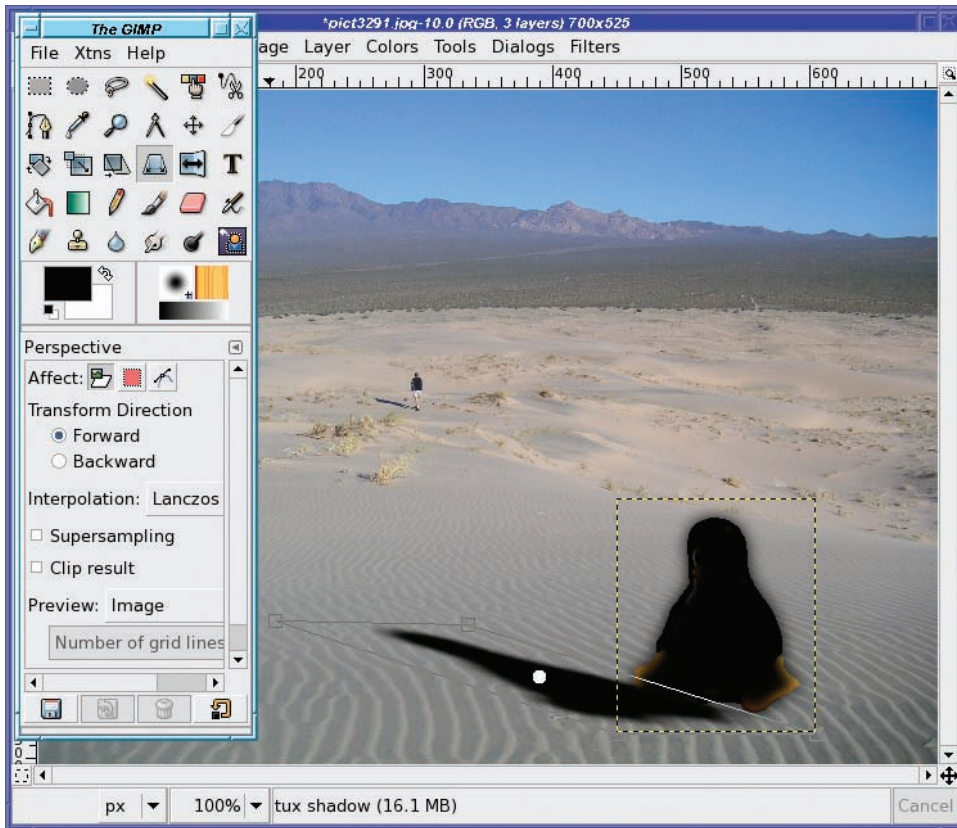


Tux is cute and all, but he just doesn't look . . . realistic. The reason is the lack of a shadow. (Well, okay, maybe it's not the *only* reason.)

## Using an Object to Cast Its Own Shadow

Start by making a drop shadow from the Tux layer.


Now use the Perspective tool to drag the shadow off in the right direction (Figure 9-24). Try to match the shadow angle that you see on other objects in the same photo. Don't worry if the base of the shadow doesn't quite match the base of the pasted object; you can fix that separately.



**Figure 9-24.** Use the Perspective tool to drag two corners of the shadow box in the right direction.

Until now, the shadow layer has been on top of the Tux layer. Obviously that's not right. Now is a good time to move the shadow layer down in the layer stack. Then use the Move tool to move the shadow so that its base is in the right place (Figure 9-25).





**Figure 9-25.** *With the shadow layer under the Tux layer, move the shadow to the right place.*

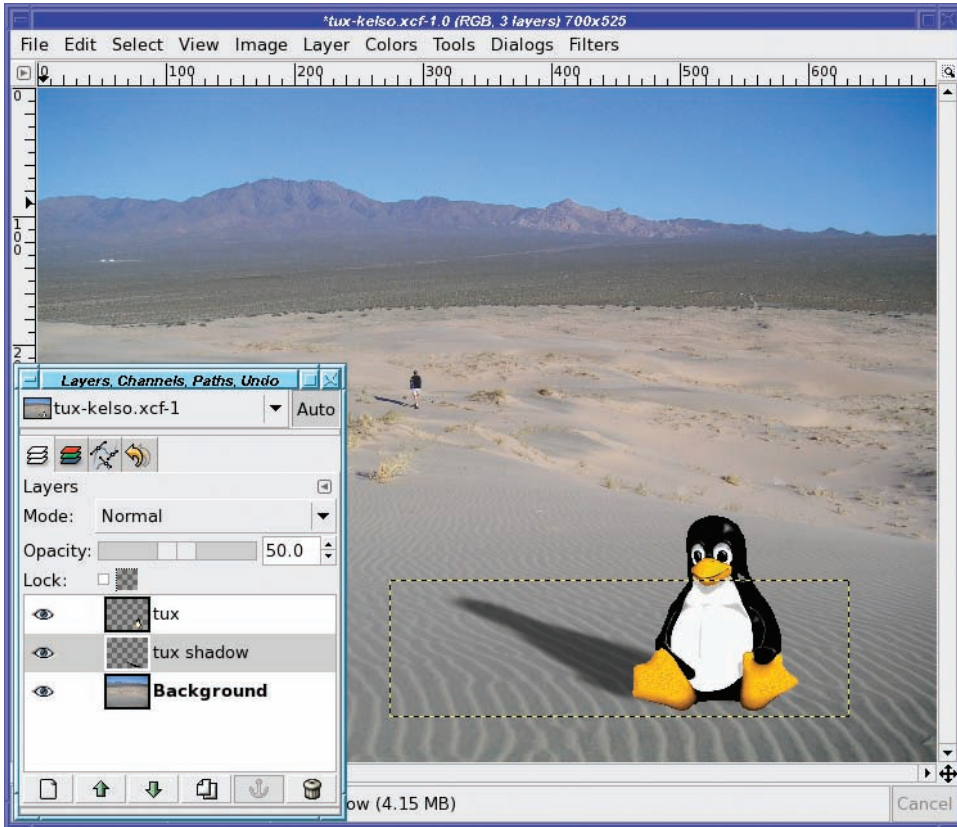
---

**Note** This is basically what the Perspective Shadow plug-in does. However, it can be hard to get the plug-in to work right; it's usually easier just to make a perspective shadow by hand.

---

## Transparency: Add the Final Tweak

That looks pretty good—but you're not quite done yet. Change the transparency of the shadow layer so that it's not completely black. Then you can see some of the texture and color underneath it (Figure 9-26).



**Figure 9-26.** Give the shadow layer some transparency, so you can see the colors and textures underneath.

Of course, sometimes you may need more work to make a shadow look perfect, especially if it's being cast onto a curved or irregular surface. But for most of your shadow needs, the Perspective tool combined with Blur and layer transparency will give you everything you need.

## Realism and Multipoint Perspective

There's a little fib in the last section. Remember where it said to match the shadow angle of other objects in the photo? That's close to true—but sometimes not close enough.


The problem is perspective (the name of the tool you used to warp Tux's shadow to the right shape and angle!).

If you look carefully, you can see the problem. In Figure 9-25, both the hiker's and Tux's shadows are at roughly the same angle. If you have a sharp eye, Tux's shadow will look just a bit out of line—tilted down. However, in the final Figure 9-26, his shadow is rotated up a bit—and somehow that looks more natural.

Why is that? It's actually a nearly ideal representation of the effect of perspective. But it's easier to start any discussion of perspective with the most elementary version.

## Single-Point Perspective

You probably remember your introduction to perspective in Chapter 4, when you drew the box for the tree planter. That was a reasonable introduction to *single-point* perspective (Figure 9-27).



**Figure 9-27.** A classic example of single-point perspective: three boxes drawn with one vanishing point

The “point” in question is the *vanishing point*. This is the place where parallel lines seem to converge in the distance. Imagine looking down the lane stripes on a long stretch of road. You’ll note after some distance you can no longer distinguish them as separate stripes. If you can see far enough, the entire road disappears to a point in the horizon.

Since there are three boxes in Figure 9-27, you can see what happens when foreground objects are offset to one side or another. But that tells you something else: the vanishing point does not have to be in the center of the image.


Use any object (such as your hand, or the opposite page of the book) to block out two boxes. Now you have an example of a drawing where the vanishing point is offset to one side.

This is similar to the effect at work in the Tux image. But in Figure 9-26, the point where the shadows converge is *way* off to one side. Why is that, and how do you figure out where the vanishing point should be?

## Two-Point Perspective


In Figure 9-27, all the horizontal lines are parallel to each other, and so are the vertical lines. The only lines that have a different angle are lines that recede “into the paper.” This is a rudimentary illusion, but we can do better.

How? By adding another vanishing point, as in Figure 9-28.



**Figure 9-28.** Perspective using two vanishing points. Note the guide layout is very different from single-point drawings.

Not coincidentally, the box has been drawn so that the angle from it to the left vanishing point is similar to the angle the shadows follow in Figure 9-26. What happens if you copy part of the two-point drawing and paste it onto the Tux picture (as in Figure 9-29)?



**Figure 9-29.** Centered on Tux (the focus of the image), the two-point cube's perspective lines show the relationship of Tux's shadow to the hiker's.

An understanding of perspective will help not only with drawing, but any time you have to make an alteration to an existing image. Getting the angles, warps, and shapes “just right” can seem impossible using LARTM (“Looks About Right To Me”), but is easy when you know how to draw perspective lines.

The GIMP doesn’t let you set up guides that aren’t exactly vertical or horizontal to help with the Perspective tool. But that’s okay: by drawing lines in a separate layer, as shown here, you can still find your vanishing points, and then let them guide you when using the Perspective tool. Then make the layer invisible before you save.

In many cases (such as this one), the vanishing point is so distant, it’s off the edge of the picture. That’s perfectly normal. Sometimes old-world draftsmen (circa 1990) would stick a pin into their drafting board way off the work, and line up their guides with a ruler.

Not so easy with a screen! But you can use the same trick you see in Figure 9-29. Make your guides smaller on a separate layer, and then scale them up to the proper size.

By using vanishing points with the GIMP’s Perspective tool and adding a realistic shadow, you can convincingly stick just about any object into even the most distorted background.

## Adding Reflections and Shading

To draw Tux’s shadow on the dunes, you had to have an idea of where the light was coming from. If you wanted to make an object look really three-dimensional, you’d want more than just a shadow: you’d want highlighting on the side of the object that faces the light, and shading on the other side of the object.

Remember shading the box in Chapter 4? This is similar, but with curved objects there are a few other details involved.

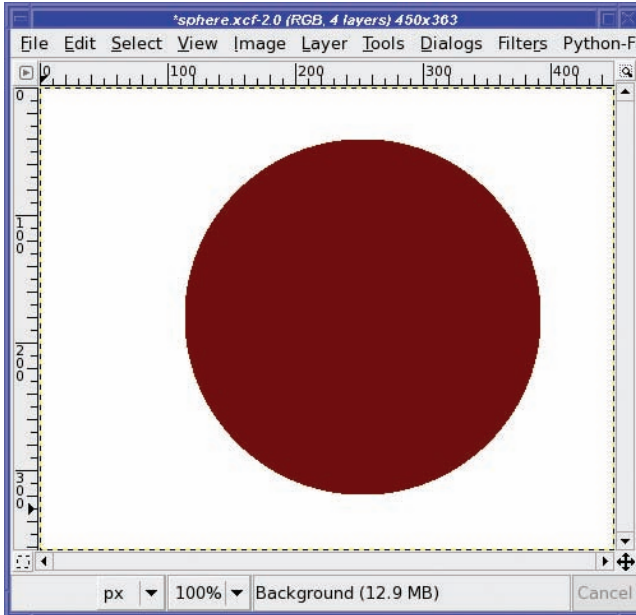
Rather than using Tux, I’ll use a simpler example: I’ll show you how to make a sphere.

---

**Note** GIMP has a plug-in to make a sphere for you, in the *Xtns* menu of the Toolbox. But once you know how to shade a sphere, you can use similar techniques to shade any object. In Chapter 11, you’ll learn more about how the Sphere plug-in works.

---

Start by making a colored circle. Make a circular selection, and then fill it using *Fill with FG Color* or the Bucket Fill tool (Figure 9-30). Use a fairly dark color: this will be your shadow color. Of course, put your circle in its own layer.



**Figure 9-30.** Start with a colored circle.

Bring up the color chooser and make the foreground color quite a bit brighter (increase its Value by sliding the V slider to the right). This will become the lighted part of the sphere.

You'll need a selection that's a bit bigger than the current circle, and feathered quite a lot. Your choice: either enable *Feather edges* in the Ellipse Select tool and make a new circular selection, or use the existing selection, enlarge it with *Select > Grow...*, and then feather it with *Select > Feather...* For my relatively small example, I'll *Grow* by 20 pixels, and then *Feather* by 45.

Move the selection toward your light source: in this case, I'll move up and to the right (Figure 9-31). Remember, you can move a selection with Alt-drag, Shift-Alt-drag, or by using the *Transform selection* tab in the Move tool options.

Once the selection is moved, fill it with your new brighter foreground color (Figure 9-32). Then change the layer mode to a mode that looks good. Screen and Dodge are good for bright colors; Soft light gives a more subtle and darker effect.



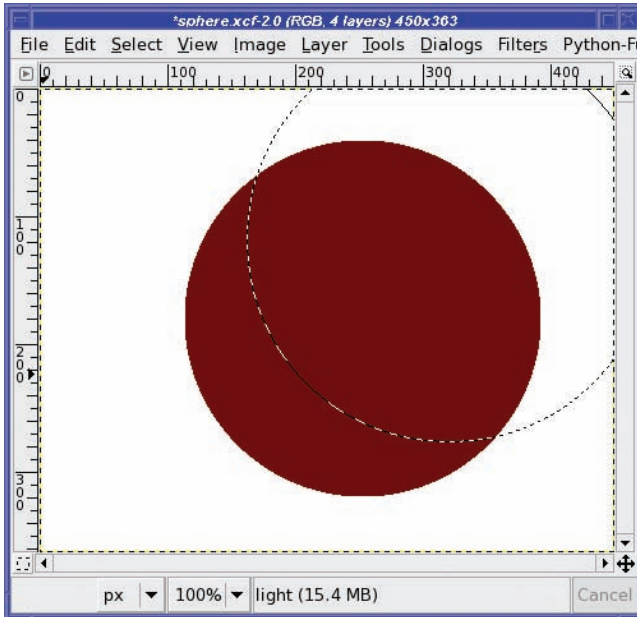


Figure 9-31. Move the selection boundary toward the light source.





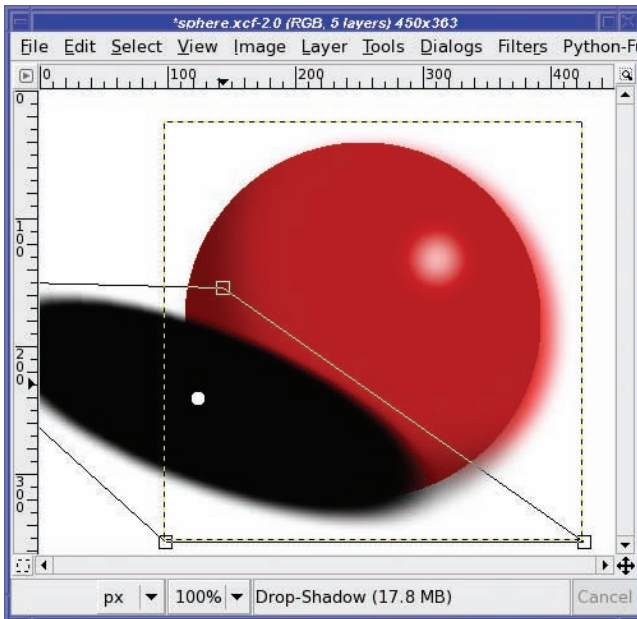
Figure 9-32. Fill with the lighter color, using an appropriate layer mode.

Add a bright highlight: make a small, feathered circular selection up near the light source, and fill it with white (Figure 9-33).




**Figure 9-33.** Fill a small circular region with white to add a highlight.

Finally, add a shadow, the same way you did with Tux. There are two slightly tricky aspects. First, if your perspective drag handles end up off the screen, it can sometimes be difficult to get them back. Second, it's hard to tell from the shadow's shape when the perspective is right, so look at the box drawn in the Perspective tool for clues (Figure 9-34).



**Figure 9-34.** *Add a shadow.*

Drag the shadow to where it looks right, and you're all set (Figure 9-35).



**Figure 9-35.** *The finished sphere*

Of course, you can modify this technique in all sorts of ways. When you need to shade a complex object, you can draw shadows along one edge of the object using the Paintbrush tool with a fuzzy brush (in a separate layer, of course), or fill complex selections you've made with tools such as the Lasso tool.

You can make spheres with different textures such as glassy, metallic, or leathery: vary the feathering of the selections you fill, and then add textures or reflected objects, perhaps warped with one or more of the Map plug-ins discussed in Chapter 7.

You can add color to the shadow, or to the highlighted areas, to emphasize color being reflected from other parts of your artwork. Try different approaches! The more you experiment with light and shadow, the better your art will look.

## Making Brushes, Patterns, and Gradients

The GIMP has a great collection of brushes, gradients, and patterns. But sometimes you need something really specific. How do you make your own?

### Making Brushes

In Chapter 4, you learned how to create parametric brushes. But what if you want to create a brush of a different type?

#### Simple Image Brushes

Any image can become a brush. But good brushes will usually be small and have transparent backgrounds.

Hey—that describes Wilber! How would you go about making Wilber (Figure 9-36) into a GIMP brush?



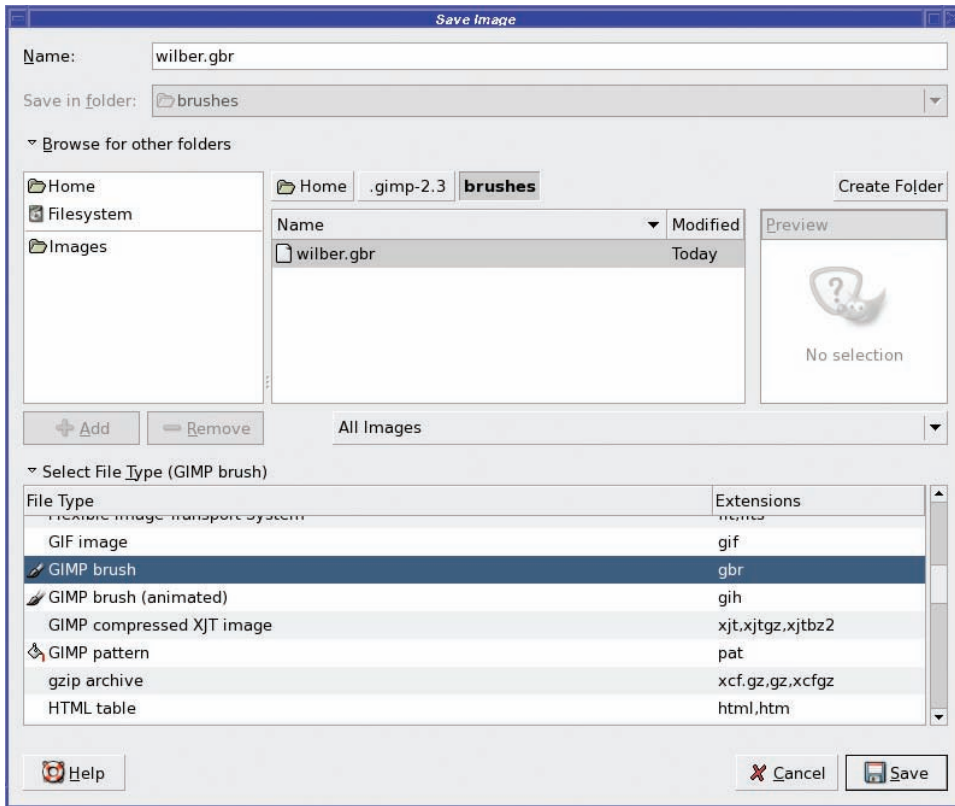
**Figure 9-36.** *Wilber, the GIMP mascot*

Making an image into a brush is easy—just save it in the right place using the right image type (Figure 9-37). The right place is the folder check-marked *Writable* in *Preferences* ► *Folders* ► *Brushes*.

---

**Tip** If you have trouble seeing your GIMP profile folder in the Save as dialog, try clicking on *Browse for other folders*; then right-click in the file list and select *Show Hidden Files*.


---



**Figure 9-37.** Save as “GIMP brush” (.gbr) to the brushes directory in your GIMP profile.


The right type is GBR, for “GIMP brush.” You can expand GIMP’s *Select File Type* list and click on *GIMP brush*, or just save it to a file ending in .gbr and let GIMP figure out the type.

Once you click *Save* in the dialog, you’ll be prompted for a *Spacing* and *Description* (Figure 9-38).



**Figure 9-38.** Options when saving a brush

*Spacing* is how far apart the images will be when you draw lines and curves with the brush (Figure 9-39).



**Figure 9-39.** *Spacing controls how far apart the images will be drawn.*

After you've saved a brush, you need to *Refresh brushes*—the button in the lower-right of the Brushes dialog—before GIMP will see the new brush. Of course, exiting the GIMP and starting again also works. Then you can use your brush.

Although it's fun to draw silly lines made of Wilbers, a more practical use for simple custom brushes is as a *stamp*, or an image you need to use frequently. Suppose you have a signature that you like to add to your artwork. You could keep your signature in a file, and open it every time you needed to use it, copy it, and then paste it into the new image. But if it's already loaded as a GIMP brush, it's always accessible through the Brushes dialog. Much easier, and you don't need to remember where you stored the file.

## Animated Brushes


The real power of custom brushes is in *animated brushes*. These are made from images with multiple layers, just like the animated image you made at the end of Chapter 3. As you drag the mouse across the canvas, instead of seeing the same object repeated many times, you'll see slightly different objects. GIMP's Vine brush, and others such as the Pencil Sketch brushes, Felt Pen, Confetti, and Galaxy, all use animation to give you a random-looking, constantly changing line. The animated brushes have a red corner in the Brushes dialog, so you can tell them from other types of brushes.

To illustrate, I'll demonstrate how to make a very simple animated brush: grass.

Start with a small image. I'll use 50 × 50 pixels. It shouldn't have any background, just a completely transparent image. (If you start with an image that has a background, *Add Alpha Channel* followed by *Clear* can fix that.)

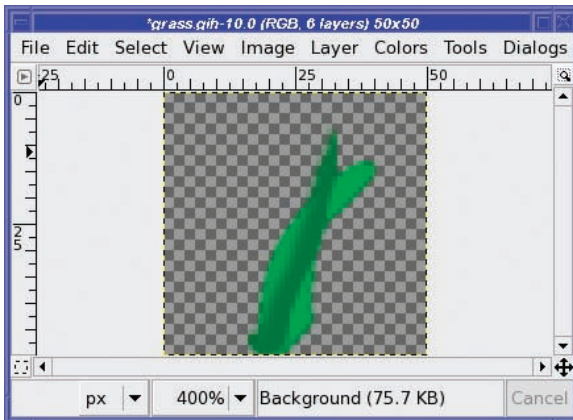
Now draw a blade of grass (Figure 9-40). Notice that I zoomed in the window to 400%: it's very hard to draw in a 50 × 50 window at normal size.





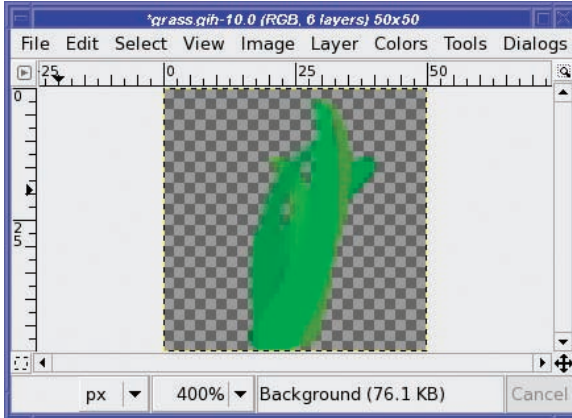
**Figure 9-40.** A single blade of grass against a transparent background

Make a new layer, and draw another blade of grass. Grass will look more realistic if each blade is a slightly different shade of green as well as a different shape (Figure 9-41).



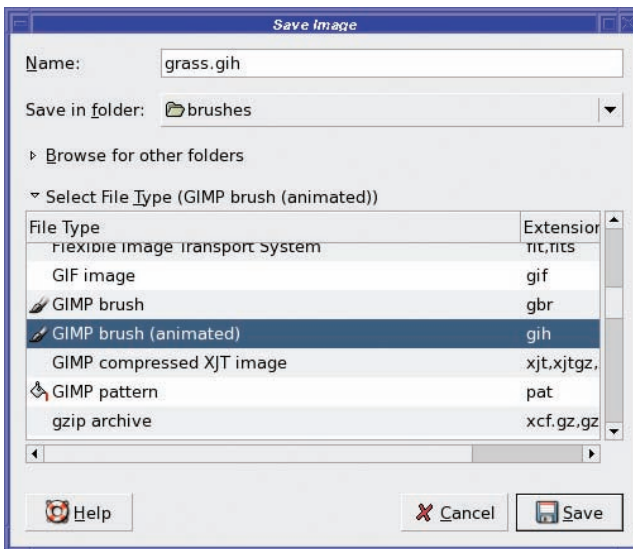
**Figure 9-41.** A second blade with a different color and shape from the first

Continue adding new layers until you're happy. I'll stop at six layers (Figure 9-42).



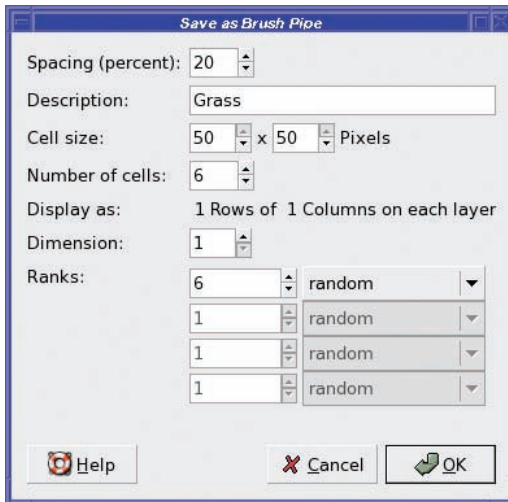
**Figure 9-42.** Six blades of grass, each in a separate layer

Now it's time to save the brush. Use the same folder you used for simple brushes: the brushes folder in your GIMP profile directory. Use the *GIMP brush (animated)* type, a *.gih* extension (Figure 9-43).



**Figure 9-43.** Save with a *.gih* extension: type “GIMP brush (animated).”

This brings up the *Save as Brush Pipe* dialog (Figure 9-44). *Pipe* is another term for an animated brush.



**Figure 9-44.** Options when saving an animated brush

*Spacing* is the distance between consecutive brush marks, just as it was with simple brushes. It's specified as a percentage of the brush size. You'll have to experiment with this setting to find the right one for your brush. For grass, 8 or 10 works fairly well.

*Description* is simply a text description you can give to your brush. *Number of cells* is the number of layers in your image. GIMP should initialize this for you.

The rest of the options—*Cell size*, *Display as*, *Dimension*, and *Ranks*—control how the frames of the animated brush are chosen from the image, and in what order. For a simple animated brush like this, leave *Dimension* at 1 and set the first *Ranks* column to the number of cells (6, in this case).

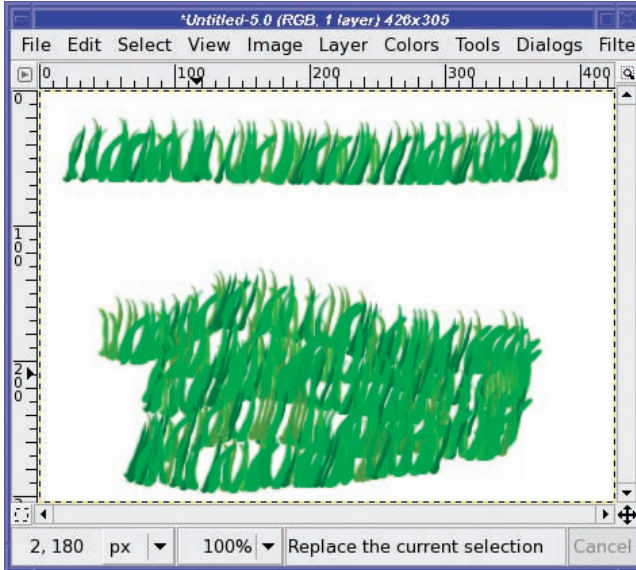
---

**Note** For more complicated brushes, you can choose to divide each layer into a grid, and then use each grid cell as a frame in the brush. *Dimension* and the other *Ranks* fields tell GIMP how the grid is set up and in what order to choose the frames. These make brush design more complicated; you're probably best off ignoring them until you get very comfortable with animated brush design. You can use them for effects like changing orientation according to the direction you're drawing.

---

The menu next to the *Ranks* entry controls the order in which cells are chosen. *Random* picks a random layer each time; *Incremental* uses the layers in order; *angular* and *velocity* let you create more advanced effects, such as changing the brush image according to the direction or speed the mouse is moving. The other options, *pressure*, *xtilt*, and *ytilt* use those values from a drawing tablet if you have one. If you don't, you're better off sticking to *Random* or *Incremental*.

Figure 9-45 shows the effect of drawing with the grass brush.



**Figure 9-45.** Drawing with the animated grass brush

## Importing Animated Brushes

The GIMP can also import animated brushes from certain other programs. Paintshop Pro calls them *tubes* and uses the .tub extension. Convert them to GIMP animated brushes by opening the .tub file as an image, and then saving as GIH. For other programs, you may be able to find plug-ins or external programs to convert a brush to GIMP's GIH format.


---

**Note** You'll sometimes see the term *image hose* to refer to animated brushes, especially in the Paintshop Pro world. The idea is that the brush is a hose that sprays out a stream of constantly changing images. *Pipes* (as in GIMP's dialog) and *tubes* are other ways of referring to animated brushes.

---

## Making Patterns


You've probably used some of the GIMP's built-in patterns by now. But you can get a pattern from anywhere. For instance, what if you want a pattern made from ocean waves (Figure 9-46)?



**Figure 9-46.** Ocean waves to make into a pattern

After making brushes, I'm sure you can guess how to save a pattern. Go to the brushes folder in your GIMP profile, and save it as type "GIMP pattern" (a .pat extension). When it prompts you for a description, type something appropriate (I used "Ocean waves"). Then open the Patterns dialog and click the *Refresh* button in the lower-right. Easy, right?

Well, of course there's a catch. When you use that pattern, you'll almost certainly see a *tiling* effect, as in Figure 9-47.



**Figure 9-47.** You can see seams at the edges of the tiles.

What went wrong? There's a sharp boundary where the bottom edge of one copy of the pattern doesn't match the top edge of the next copy (and similarly with the left and right edges, though the effect is less obvious). The eye is very good at picking out boundaries. That won't do!

Fortunately, the GIMP has a filter to help with creating patterns which can be found at *Filters* ► *Map* ► *Make Seamless*.


After running *Make Seamless* on the waves pattern, saving it, and reloading it in the Patterns dialog, now *Fill with Pattern* gives Figure 9-48.



**Figure 9-48.** After running *Make Seamless*, you can't see the boundaries any more.


There are also plug-ins you can download, such as Resynthesizer, and other techniques such as a high-pass filter and the Clone tool, which can go beyond Make Seamless to help make your patterns tileable.

Running Make Seamless sounds easy, and it is. But if you make patterns often, you'll probably find that it's not always that simple. Regular patterns, such as bricks, may end up warped or blurred in unnatural-looking ways. Even fairly random patterns like waves don't always work. For instance, my ocean sample was a small area selected out of a larger image. But I get a much worse result if I move the selection just slightly over, to a place where there's a dark spot visible beneath the surface (Figure 9-49).



**Figure 9-49.** Another selection from the same image: ocean waves with a visible dark spot

Even after Make Seamless, using this image as a pattern shows a tiled effect (Figure 9-50).



**Figure 9-50.** This pattern doesn't look realistic when it is tiled.

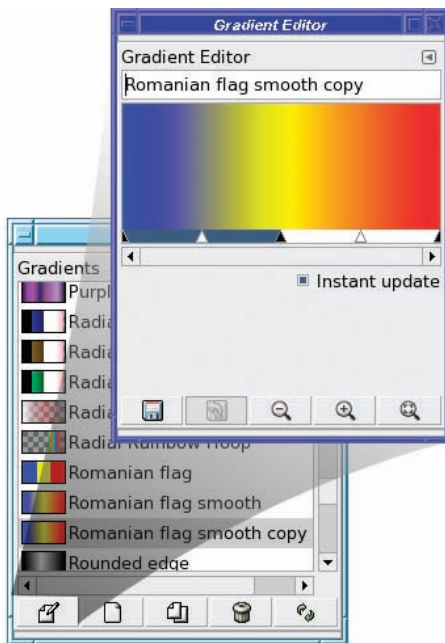
The key to choosing a good pattern is finding one without much change from one edge to the other. If it's brighter in one corner, if there's a spot in the middle, if there's an obvious shape inside . . . it will be obvious where one tile ends and the next begins. Sometimes this is acceptable (and even some of the GIMP's built-in patterns show obviously tiled edges), but most of the time you won't want that.



## Making Gradients

You can make your own gradients, too, by using the *Gradient Editor*.

The easiest way to learn the Gradient Editor is to edit an existing gradient. Select a gradient similar to your goal, duplicate it, then click *Edit* to get the Gradient Editor (Figure 9-51). You can also find it in the context menus of the Gradients dialog (right-click on any gradient in the list, and then select *Edit* to edit that gradient, or *New* to make a new one).



**Figure 9-51.** *The Gradient Editor*

---

**Caution** Make sure you duplicate an existing gradient before running the Gradient Editor. If you try to edit one of GIMP's built-in gradients, many functions in the dialog won't work because the gradient will be opened as "read only."

---

At the top of the Gradient Editor is the name (for a new gradient, "Untitled"). You can change this to the name of your new gradient.

Below the name is an image of the gradient as it currently looks, and below that, an area for *range selection sliders* (the blue-and-white bar with black-and-white triangles).

In the range selection area, black triangles define *segments*. Each segment has two endpoint colors. White triangles represent the midpoints of ranges. The currently selected range is blue; the others are white.

In Figure 9-51, there are two segments. The first (leftmost) segment is active, since that's the one with the blue bar below it. Its left endpoint color is blue, and its right endpoint color is yellow. The second segment has a yellow left endpoint, and a blue right endpoint.

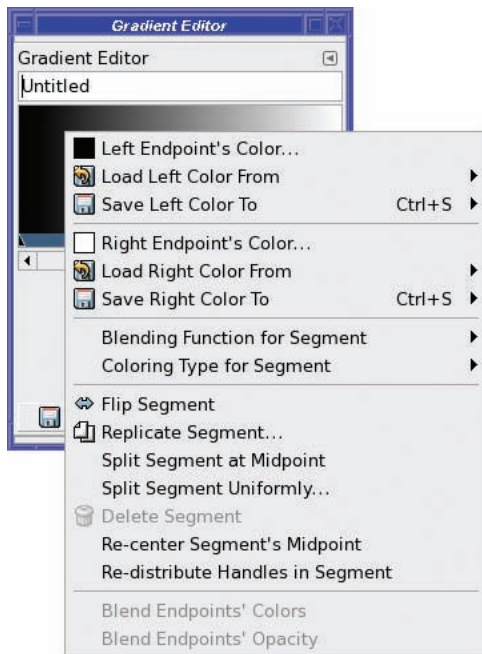
Each segment will fade from the color at one endpoint to the other. How fast it will fade is controlled by dragging the midpoint slider (the white triangle). Drag it to the right, and you'll see more blue; to the left, and you'll see more yellow.

---

**Tip** To make an abrupt cutoff, like “Romanian Flag” (as opposed to “Romanian Flag Smooth”), make the left and right endpoints the same colors.

---

You can control many aspects of blending, and add or delete segments, by right-clicking in the gradient to get a context menu (Figure 9-52).



**Figure 9-52.** *The Gradient Editor's context menu*

When you've edited a gradient to your liking, the *Save* button adds the gradient to your GIMP profile. You don't need to choose a filename or file type for it: GIMP will handle that for you.

Most people won't ever need to create or modify a gradient. But for those who do, GIMP's Gradient Editor gives you plenty of power to do whatever you need. Experiment and see what it can do for you.

## Summary

By now you know how to create your own brushes, patterns, and gradients. You know all about perspective and how to use it when drawing, how to make shadows, and how to make text or other objects fade out using layer masks. And you have some idea of the effects you can achieve by combining the GIMP's many layer modes in your drawings.

Next, it's time to revisit photographs, and see how you can use some of these same techniques to make your photos better, or to combine several photographs together . . . a process called compositing.





# Advanced Compositing

**C**ompositing is a fairly general word. In its most general sense, it means combining two or more images of any sort—even small elements such as a circle, a shadow, or a text layer—to make a new image. You’ve already been doing that throughout most of this book.

In this chapter, I’ll cover a more specific meaning: blending two or more existing photographs to make a new image. (In some cases, the two photographs may be copies of the same photograph, as you’ll see.)

You’ll learn some tips for improving photos that weren’t covered in the chapters on photo retouching, and tips for extracting detail from a photo that you may not have known was there. You’ll learn to create artistic renderings of your photos (without needing GIMPpressionist or similar plug-ins), and how to combine photos in interesting ways. You’ll learn how astronomical photographs are made, and how to use astrophotography techniques to take photos you didn’t think your camera could capture. Finally, you’ll learn how to make panoramic images, including some tips on how to shoot them in the first place. The chapter will cover the following:

- Colorizing images
- Improving photos by self-compositing with layer modes
- Making photos into art: self-compositing with modifications
- Compositing unrelated images
- Image stacking
- Stitching panoramas

## Colorizing Images

Did you ever have an old black-and-white photograph you wished was in color? Or maybe you have a color photograph, but you don’t like the real colors. You may want to colorize a whole image, or perhaps just a part.

*Colorizing* a black-and-white image is really just a variation on the sepia photo technique you learned in Chapter 8. Instead of using a light brown tone, just experiment with various colors until you hit the combination you’re after.

But what if you want to alter the colors in an image that already *has* color?

For instance, have you ever wondered what your car would look like in another color? A professional paint job costs a fortune, but with the GIMP it's both cheap and fun (Figure 10-1).



**Figure 10-1.** The car is okay in blue, but would it look better in red?

Begin by selecting the part of your photo you want to colorize using the techniques you learned in Chapter 5. In this case, that would be all parts of the car that are blue.

---


**Tip** For this kind of operation, you'll often find that a combination of selection techniques works best. On the car, it was easiest to start with Fuzzy Select (Select Contiguous Regions), which caused some overrun. Then use the Lasso or the QuickMask to fill in selection gaps and cut away any excess.

---

Copy your selection and paste it, which will make a new floating layer. As usual, convert it to a regular layer with the *New Layer* button (at lower-left in the Layers dialog). If you made the background layer invisible, your new layer might look like Figure 10-2.

Now, color the layer. You can use the *Colors* ► *Colorify* filter (Figure 10-3). Look for *Colorify* in the *Filters* menu in GIMP versions prior to 2.4. Or just use *Fill With FG Color* to get a flat color—it won't look good right away, but it will look fine when you combine it with layer masks.



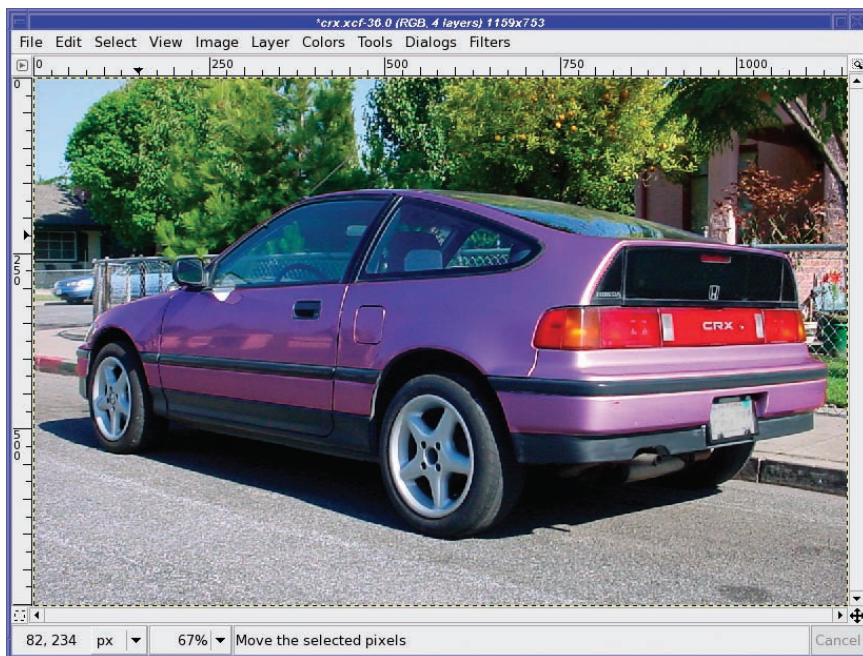


**Figure 10-2.** Make a layer consisting of just the colored part of the car.



**Figure 10-3.** Now that looks racier!

But don't stop there. As long as your color layer is different from the original, you can apply layer modes to see yet another array of effects—for instance, a red layer in Overlay mode turns the car purple while keeping the metallic look of the original blue paint (Figure 10-4).



**Figure 10-4.** *Overlay mode turns the car purple.*

Of course, you can make several copies of each part of the car and color them differently, or just paint solid colors on the overlay layer; you can make a two-tone, or add stripes. Then switch layers on and off to see which color or pattern you like best. (This is also a good way of testing hair colors on a photo of yourself before you take the plunge and dye your hair magenta!)

Some modes, such as Multiply, will probably look strange at first. Don't forget to also try changing *Opacity*, which will usually make any mode give a workable result. It's very difficult to predict the effects of the layer modes on various color shifts you may make, even if you're experienced with these modes. The bottom line? Try them all and see.

## Improving Photos by Self-Compositing with Layer Modes

Think way back to Chapter 2, where you used tools such as Levels and Curves to improve the brightness of an image.

You can make similar changes by using layer modes, sometimes with a much better result. Just make a copy of the image, right on top of the original, and experiment with the layer mode of the top copy.

Let's take a look at some of the corrections you can make by compositing an image with itself.

## Using Screen Mode for Dark Images

Sometimes you pick the wrong f-stop. Sometimes the camera guesses the exposure wrong. Especially aggravating, sometimes there just isn't enough light where you want to shoot—such as in Figure 10-5.



**Figure 10-5.** A dark indoor shot from Las Vegas

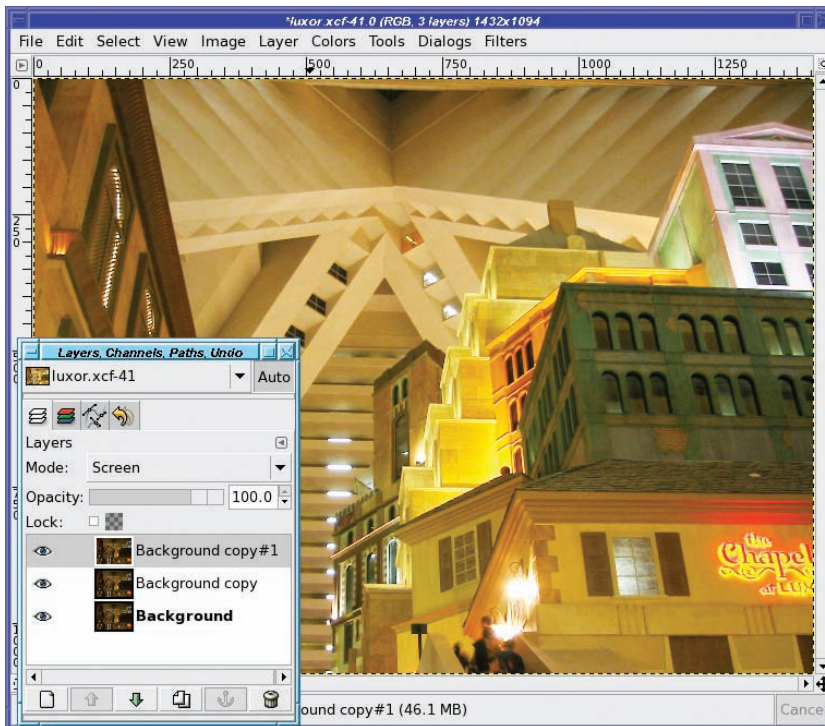
When your picture comes out too dark, you can try to adjust the color, brightness, and contrast to get the perfect balance. But there's a much quicker way to get nearly ideal brightening.

First, go to the Layers dialog and duplicate the Background layer. Next, change the resultant layer to Screen mode. Usually, that's all you have to do.

In some cases, the result will be too light, but you can adjust that by turning down the *Opacity* slider. Other times, such as with this picture, the original was so dark you need to lighten it further. In that case, just duplicate the second layer and you'll get a third layer, already in Screen mode. The picture will become even brighter. Figure 10-6 is the result of two additional screen layers.

Of course, if your second or third layer goes too far, you can adjust *Opacity* on just that layer to get exactly the right level.





**Figure 10-6.** Look at all the detail you can see in areas that were previously just mud, particularly the building at center right.

---

**Note** In some cases, especially when you're just trying to lighten up a portion of a photo, you'll find Dodge mode works a little better. But most of the time, Screen is just right.

---

## Using Overlay or Hard Light When Light is Flat

When you find yourself confronted by an image that's just way too flat and has no significant bright areas, self-compositing using Overlay or Hard light mode can help bring out color and detail.


Hard light will tend to overwhelm any bright areas in the photo, so confine it to dull, washed-out images only. Figure 10-7 shows the striking difference it can make to a dull image. For images with very bright colors or a lot of detail, Overlay is more versatile.



**Figure 10-7.** *Hard light is an easy and fast way to bring out color in an image that has no bright areas. The left half of the image is the original; the right half has been self-composited with Hard light.*

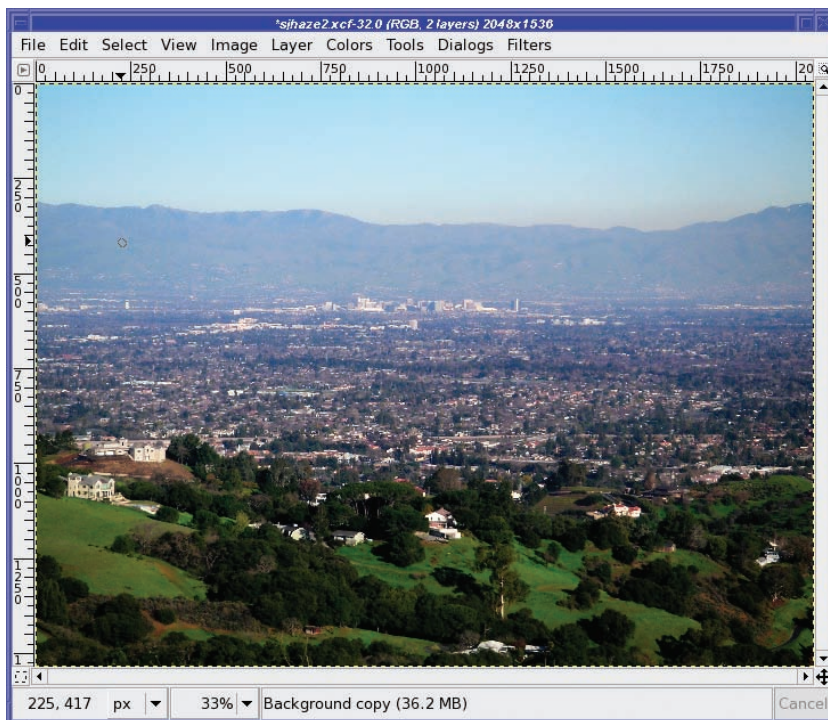
## Using Overlay or Burn to Cut Through Haze

I live near a major city, and it's frequently smoggy or hazy. When I try to take long-distance shots with my camera, I'm often disappointed: objects that seemed perfectly clear to the eye are buried in haze in the photograph, as in Figure 10-8.



**Figure 10-8.** *You can barely see the city for all the haze.*

Overlay mode excels at cutting through haze. Make a layer copy, and then set the top layer to Overlay. The city becomes much clearer (Figure 10-9).



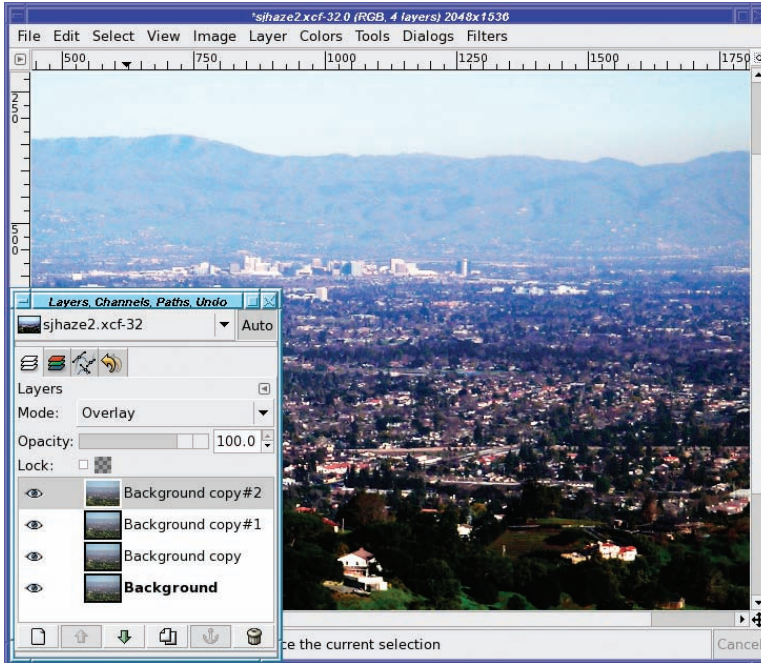
**Figure 10-9.** *Overlay mode cuts through the haze a little.*

But that's just the beginning. With the Overlay layer selected, click the *New Layer* button again. GIMP makes another copy, already in Overlay mode—and the distant city becomes sharper. Click again, and it becomes sharper still. You can keep adding Overlay layers until the image gets as sharp as you need (Figure 10-10).


If that effect isn't strong enough for you, try turning off all but the first two layers and using Burn mode on the upper layer (Figure 10-11).

Burn mode gives excellent detail and color on the city that previously was shrouded in haze. The down side is that the rest of the image becomes black and you no longer have those nice green hills in the foreground.





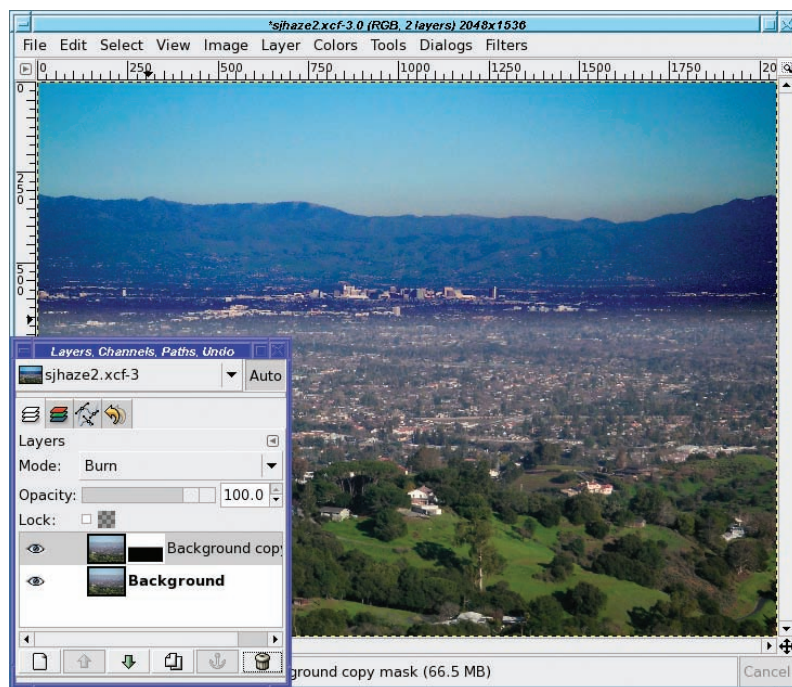
**Figure 10-10.** Stack multiple layers in Overlay mode to cut through haze even more.



**Figure 10-11.** Burn mode cuts through the haze much better than Overlay.

Of course, you can adjust the strength of the effect by changing the opacity of the upper layer. But with Burn, you may find that light areas are still too dark and you can't get the nice green hills you had when you started.

Adding a layer mask to the burned layer can fix that. The goal is to make the layer opaque where you want the effect, fading to transparent in places you don't want to change. To do that, either paint in the area you want burned, leaving the rest black, or draw a gradient separating the part you want burned from the part you don't, as in Figure 10-12.



**Figure 10-12.** A layer mask with a gradient can preserve the foreground hills while still showing detail on the distant city.

## Making Photos into Art: Self-Compositing with Modifications


As you've seen, compositing an image with itself usually affects the contrast, colors, and apparent sharpness of the image. It's a useful technique for changing iffy photos into better ones.

But if you modify one of the copies of the image before you composite them with layer modes, you can create many other interesting artistic effects.

### Making “Drawings” and Other Effects Using Layer Offset

The idea is to start with any photo you like, duplicate the Background layer, and then move it just a little bit. It's easier to describe what's going on after you see an example of the technique.

In the first series, you'll see three effects applied to the same image (Figure 10-13).



**Figure 10-13.** A drawing capturing all the petal detail could take quite a long time.

The top (duplicated) layer will be moved down and to the right by a few pixels.

---

**Tip** A good way of moving a layer by small amounts is to select the Move tool, and then use the arrow keys to move the layer one pixel at a time.

---

Each caption will tell you the exact number of pixels moved to achieve the effect shown. Your own experiments will probably use different shifts. Less movement will make finer lines; more will add detail and color. Sometimes you'll like more vertical movement, sometimes more horizontal. Fortunately, it's cheap and easy to experiment!

### Colored Pencil Drawings: Divide Mode

Generally, you'll move the image only moderately when using *Divide* mode. The first "rendering" is a fairly strong example of the effect, showing heavy shadow detail by moving the top layer *down* more than sideways (Figure 10-14).



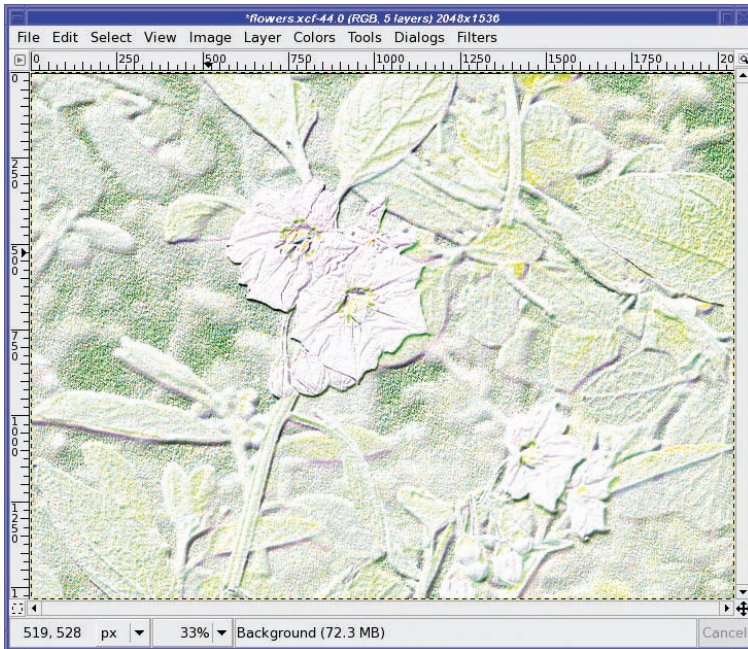


Figure 10-14. The top layer was moved three pixels right and six down in Divide mode.

It's striking that such a dark image results in such a light drawing. Also, the complexity in the dark areas has created an effect very much like you would see when an artist draws with colored pencils on heavily textured paper. Figure 10-15 shows another example, with a more complex landscape photograph.

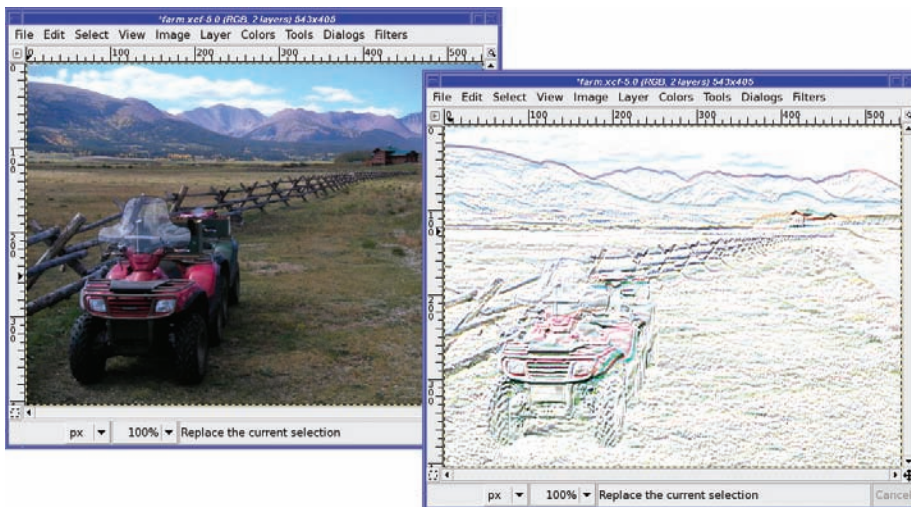
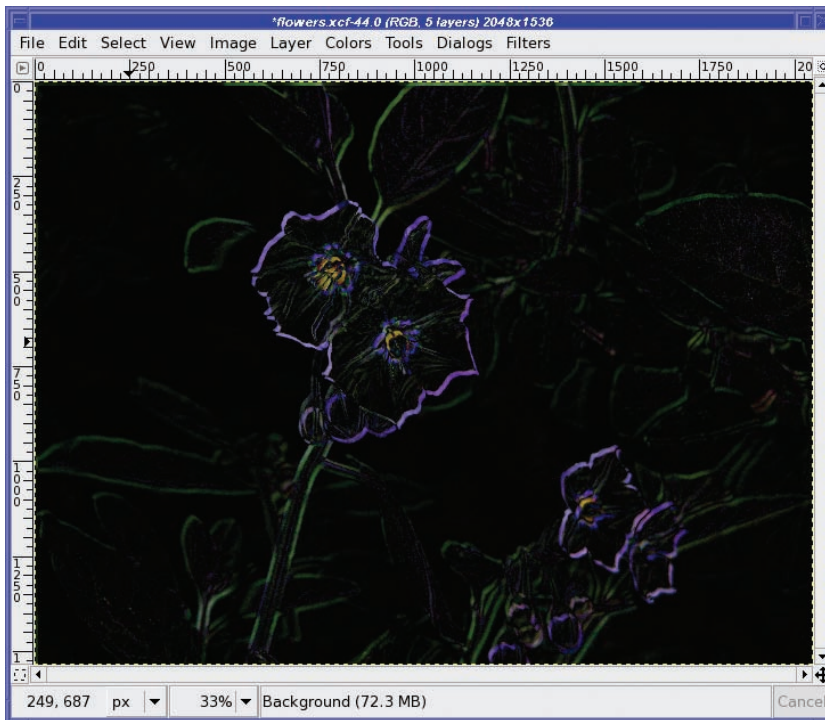


Figure 10-15. The “pencil sketch” effect on a landscape photo

Less complex images will generate fewer lines, and can be moved less to get a more sketched look. You can further modify your “drawing” by playing with the colors, or turn it into a graphite pencil sketch by converting it to grayscale.

### Black Velvet Paintings: Difference Mode


Sometimes, you want a dark result rather than something bright. You’ll almost always get a pleasant pastel image from Divide mode. To get the opposite effect, use *Difference* mode on the same picture to create Figure 10-16.



**Figure 10-16.** *Difference mode makes a neon line drawing on black. The top layer was moved eight pixels to the right and ten pixels down while in Difference mode.*

That’s a neat effect, almost like neon lights. But playing with the *Opacity* of the difference layer can make the effect more interesting, and turns the artwork into a “black velvet” painting (Figure 10-17). Too bad I don’t have any photographs of Elvis to use as examples!

Difference usually requires you to move the top layer more than in other layer mode effects. This picture is no exception, having moved only a little further right than Figure 10-17, but almost three times as much to the right. (When moving a layer over larger distances than a few pixels, remember that holding down Shift while pressing an arrow key moves the layer by a larger distance than by using the arrow key by itself.)



**Figure 10-17.** Reducing the top layer's opacity to 75% brings out some of the details within the lines.

The result will often change dramatically as you vary the offset of the two layers, sometimes in surprising ways. If you go too far, particularly in Difference mode, the “painting” will start to become blurry. It’s an interesting balance that looks like what might happen if Van Gogh had tried to imitate Rembrandt’s dark realism.

You can get a similar effect by using Subtract mode, but in most cases, Divide will give a more pleasing image. However, it’s worth a try—when Subtract does work, it’s spectacular.

If you look closely, sometimes you can see the exposed edges of the background layer at top and left. Of course, you can crop the image to eliminate those edges before saving your painting.

### Embossing: Grain Extract

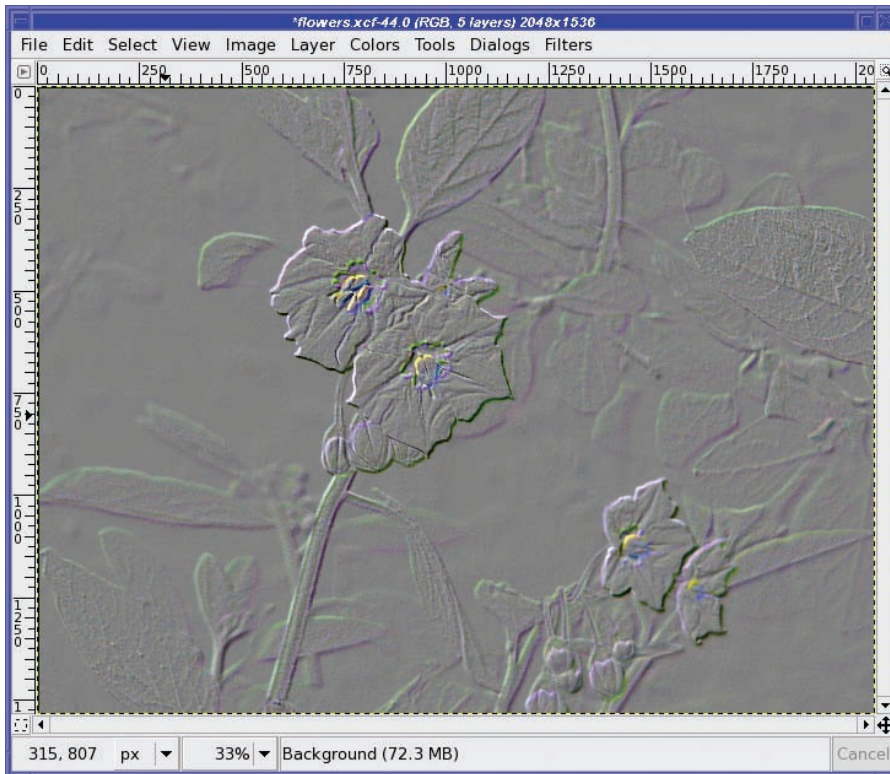
Although *Grain extract* was originally intended for storing information to be used in simulating photographic grain, it’s actually most useful for making an *embossed* version of an image. It works with almost any image, so it should certainly work with these flowers.

Normally, you’ll use the least layer movement with Grain extract (Figure 10-18).

This is a particularly strong example, partly to make sure the effect shows up well, and partly because the color accents are attractive.

In a real embossed work, the image would be colorless. Many other “emboss” conversions will use as little as one or two pixel shifts, and won’t show much color at all, if any. You can always convert to grayscale to eliminate the color, of course.





**Figure 10-18.** The top layer was moved five pixels to the right and five pixels down while in Grain extract mode.

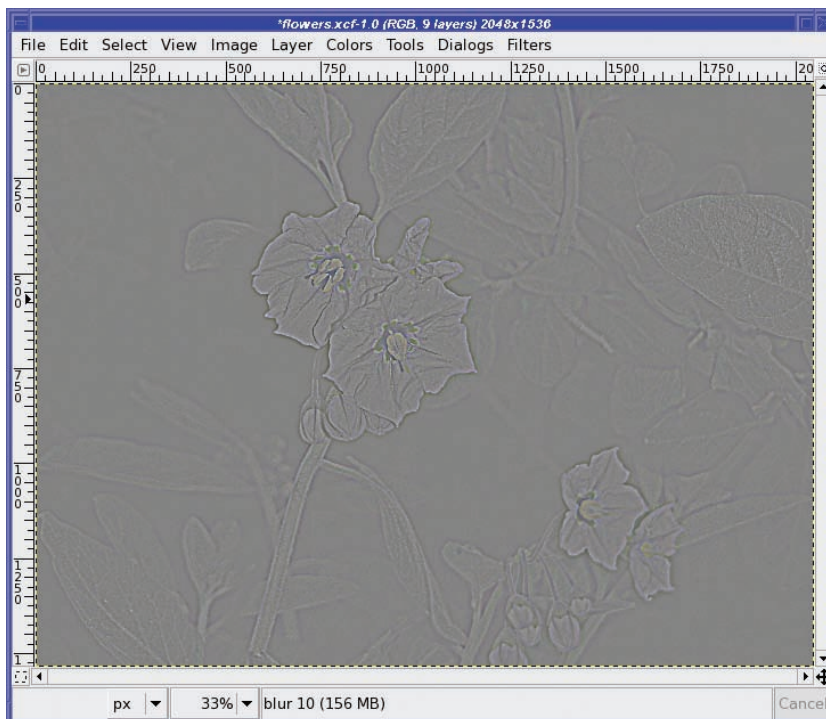
## Adding Blurs and Other Tricks for Artistic Effects

*Gaussian Blurring* the top layer instead of offsetting it before using Grain extract gives an effect like a tintype photo (Figure 10-19). Reversing the layers—putting the blurred layer beneath the original rather than above it gives a similar effect except that the colors are reversed—the flowers are darker than the background, instead of lighter.

Sometimes you can use Gaussian Blur instead of an offset when using other layer effects, like Divide. Typically, a blur gives a somewhat weaker, more subtle effect than an offset. Edges won't be emphasized as strongly.

There's a compromise between the two: use a *Motion Blur*, which blurs in a specific direction. Since it's directional, Motion Blur acts a lot like an offset; but it smooths out the effect, giving less emphasis to the edges and more to the gradations within the image. The color accents are also muted. The difference can be subtle, but try it and see whether you like it better.

By now, you probably see what's going on. Most images will have sudden changes in color where they transition from one object (a flower) to another (the background). These artistic effects all capitalize on those transitional edges. Anything that causes the edges to spread or interfere will give an interesting—and sometimes unexpected—result when you use Divide, Difference, Subtract, or Grain extract.



**Figure 10-19.** Use *Blur* and then *Grain extract* for a tintype effect.

## Compositing Unrelated Images

So far, you've used layer modes to combine an image with itself (perhaps after minor modifications such as blur or offset). But, of course, you can use layer modes to combine two different images as well.


### Using Soft Light for Combining Images

*Soft light* mode is a real workhorse for combining just about any two images (Figure 10-20). Copy an image and paste it as a new layer over a different image, and you'll probably get a fun result (Figure 10-21). As a rule, use the less distinct image as the background, since it will slightly dominate the combination. You can also use *Soft light* to give the effect of objects reflected from water or glass.

As a strange side-effect, self-compositing with *Soft light* doesn't do what you'd expect: it doesn't soften the image. Instead, you'll get a very nice increase in brightness and contrast. Even pictures you think are just fine can be surprisingly enhanced, with no fake look at all.



**Figure 10-20.** *Two unrelated images to be composited*



**Figure 10-21.** *Combine the two unrelated images with Soft light.*

## Using Overlay for Dark Images

Overlay has an effect quite similar to Soft light, and is always worth trying to see how the effects compare. Overlay is also good for adding elements of an image on top of darker areas in another image, like the photographs in Figure 10-22. You'll still see elements of both, of course—it's not a substitute for painstakingly selecting an object and pasting it into another picture.





Figure 10-22. Unrelated images: Shiprock, New Mexico, plus real sailboats

But sometimes you might not mind “ghostly images.” For instance, the ghostly sailboats in front of dark, looming Shiprock give a nice effect (Figure 10-23).

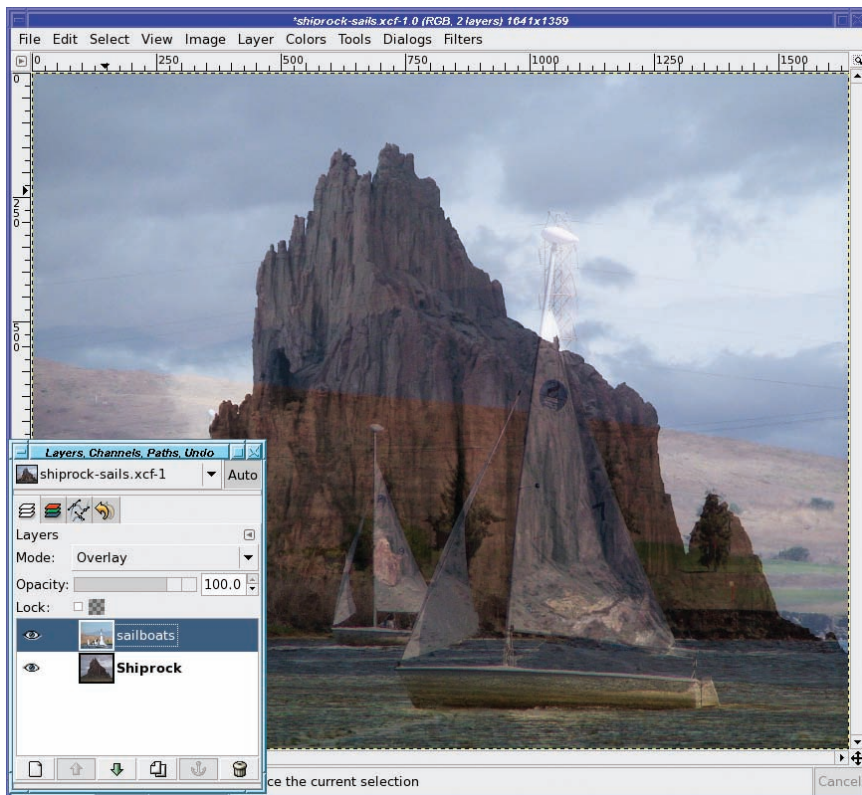
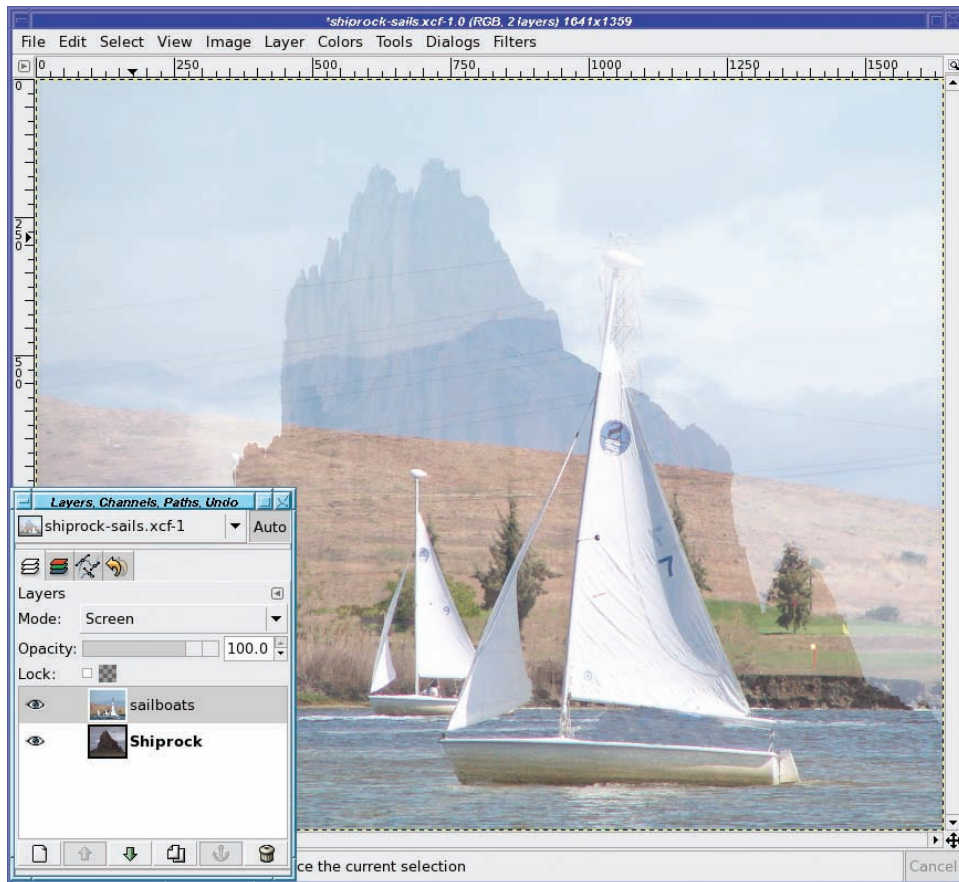


Figure 10-23. Overlay tends to make dark images, creating a moody effect.

## Using Screen to Get a Lighter Effect

*Screen*, too, is a useful tool for combining images—but it has a very different effect. *Screen* brightens the image, so instead of ghostly, looming figures, the brighter image will stand out (Figure 10-24).



**Figure 10-24.** *Sailboats in front, in Screen mode; Shiprock in back*

## Using Addition to Complement Light and Dark

*Addition* mode can be useful when you have two dark images that might work together, or two images with complementary light and dark areas (one is light where the other is dark), such as in Figure 10-25. Combining them in *Addition* mode gives the result in Figure 10-26.



Figure 10-25. It's hard to see how these could relate.

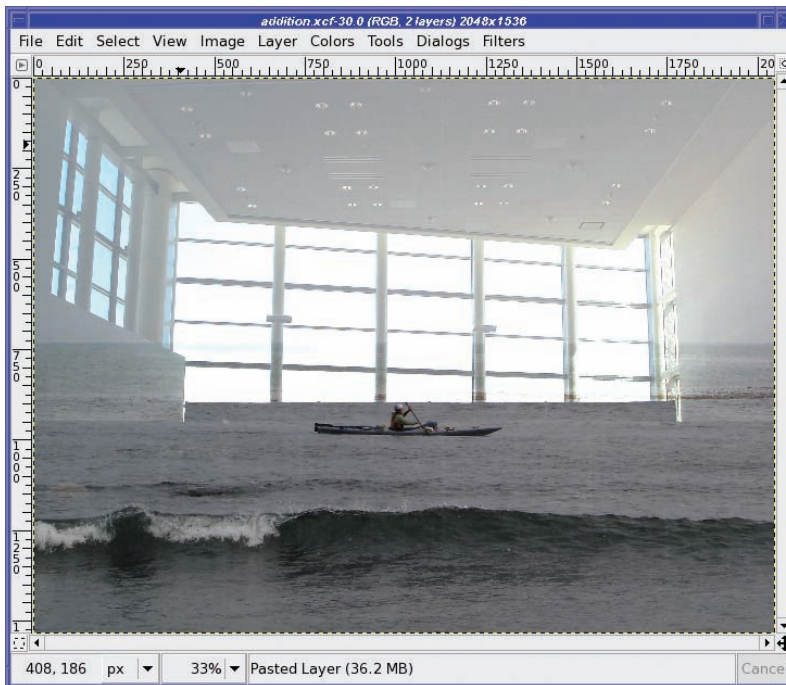


Figure 10-26. Addition mode can combine dark areas of two images.


## Using Subtract to Make a Cut-out Mask

*Subtract* usually inverts dark and light on the layer being subtracted. You can use it to take an image that has light and dark areas, such as the gull in Figure 10-27, and make a mask for another image, as shown in Figure 10-28.






**Figure 10-27.** The gull's light and dark areas can make a mask for the building.



**Figure 10-28.** Use Subtract mode to make a mask for another image.

Difference will give a slightly different and weirder effect, turning the dark colors into a different color (Figure 10-29).




**Figure 10-29.** *Difference adds color to the previously dark areas of the mask.*

## Making Eerie Colors with Burn

When the background picture has very distinct darker and lighter areas, like the tree silhouettes in Figure 10-30, it can seem to turn into a mask when Burn mode is used (Figure 10-31). In this case, the trees blot out the strange rock formation, but the lighter sky alters its colors to appear otherworldly. You've probably seen effects like this used to illustrate science-fiction book covers.



**Figure 10-30.** *Burn works best when one image has distinct light and dark areas.*



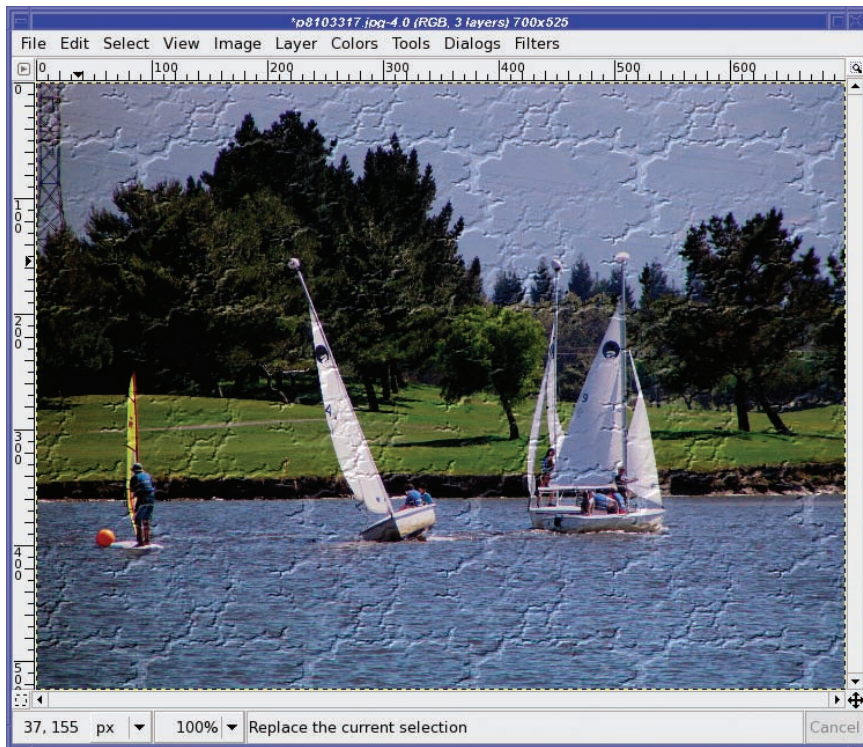
**Figure 10-31.** Burn mode can combine a sky or other light, colored area with another image for an interesting effect.

## Using Grain Merge to Add Texture

To add texture to a photograph, use the photo as the lower layer, and fill the upper layer with a tileable pattern, or noise created using Grain extract.

For instance, using the GIMP's built-in pattern named "Crack" as the upper layer in *Grain merge* mode, on top of the sailboats, gives the result shown in Figure 10-32.

As always, when dealing with layer modes, the results of any particular combination tend to be unpredictable. These examples can point you in the right general direction, but experimentation will often give surprising results. You can duplicate layers and use multiple modes with varying degrees of opacity to create almost *anything*. It's fun and easy to try several options, and it's almost always rewarding.



**Figure 10-32.** A pattern in Grain merge mode adds extra texture to an image.

## Stacking Images

Have you ever marveled at a space shot, like the ones from the Hubble Space Telescope? Did you know that the “raw” images straight from the camera are hardly ever that good?

Astrophotography is one area where image compositing is an essential tool. Professional observatories, NASA, and amateurs in their backyards all use a set of related techniques known collectively as *stacking*.

The first image stacks were just that: photographic negatives or slides stacked directly on top of each other. When these stacks were viewed on a light table, or printed, the result was a photograph with much higher contrast. With stacking, details that weren’t visible in any single original became easy to see.

Of course, today, stacking is usually done by computer. Stacking multiple images of the same object is very similar to the self-compositing you’ve already used earlier in this chapter. But compositing several different but similar images is better than self-compositing in many cases, as you’ll see.

Sometimes astronomers get carried away with stacking. The famous “Hubble Deep Field” photo captured more than 1,500 distant galaxies in a tiny patch of sky 1/30th the diameter of the full moon. The Deep Field project shot 342 images over 10 consecutive days, and used 276 of them in the end. Of course, the Hubble team uses much more complex stacking methods than simple layer modes. But many of the basic principles are the same.



Although stacking techniques were developed first for astronomy, you can stack earth-bound images as well. You can boost the light level of a dim image, improve contrast, or even increase resolution. It can also be fun to stack several different images into an “average.” But stacking will always add another benefit: reduced noise.

## Reducing Noise

Every photograph includes some unwanted interference, from thermal activity in the camera's CCD chip or various other sources. (Such as cosmic rays. Really!) Collectively, these spurious signals are called *noise*.

Noise is random: it shows up in unpredictable places. If you take two images of the same scene, the scene may not change, but the noise will be different.

Suppose you take an image that has a light speck due to noise. If you composite that image with itself, or use a tool such as Levels or Curves to boost the contrast, you will increase the speck of noise as much as the object you're trying to capture. If you've tried boosting contrast on a very dark photograph, you've probably seen that already.

But if you composite two *different* images of the same object, chances are that the second image's noise will be in different places. The noise in the first image will be partially cancelled out by the second image, and vice versa. The result is a less noisy image.

So how do you go about making a stack? There are three steps:

1. Open the images as layers.
2. Register the layers.
3. Combine the layers.

## Loading All the Images as Layers

Load each image in the stack as a separate layer. Open the first image normally, and then use *File* ► *Open as Layer...* for the rest. If you have more than two images, use Shift-click or Control-click in the file selection dialog to select all your images at once, and they'll all open as separate layers.

Once all the layers are loaded, make all but the first two layers invisible. Next, before you stack the two layers, you have to register them.

## Registering the Images Using Difference Mode

Stacking won't work well unless the images fit over each other exactly. If the subject of the photograph isn't in the same place, or isn't the same size, then it will look blurry when the two images are combined.

You can make this step much easier if you use a tripod when shooting images you plan to stack. If you try to shoot hand-held, you will almost certainly move or rotate the camera. That can be fixed, but it's tedious: use GIMP's Measure tool to measure sizes and angles in the two images, and then Scale and Rotate by the ratio of the two measurements. Eek! It's much better to use a tripod and get them right in the first place.

---

**Note** Even the best tripod will not give you perfectly registered photos when you take pictures of the night sky. Why? The earth is turning, which means, from the camera's point of view, that the sky is moving—about a degree every four minutes.

---

Once the images are approximately aligned, select the Move tool in the Toolbox and set the mode of the upper layer to Difference. This will show every place the two images differ. Two perfectly registered, identical layers will look totally black when the top layer is in Difference mode. Otherwise, move the upper layer around (remember the arrow keys for moving by one pixel at a time) to make the combined image as black as possible.

Now it's time to make a stack, using one of the methods below. After you've stacked one layer, make the next layer visible, register it, and stack it. Repeat until you're out of layers.

## Increasing Light by Additive Stacking

There are times when a flash just isn't appropriate. Did you ever go out at night with your little digital camera and wish you had a fancy SLR that could take time exposures?

With stacking, you can—even if your camera can't. Just take a lot of short exposures, and then stack them in Addition mode.

For instance, I shot some photos of a fence near a local reservoir. The camera I was using wouldn't do an exposure longer than four seconds. At that setting, the original images just look black. It's hard to tell that there's anything in the shot at all.

If I take a single image and use Curves or Levels on it, I can see that I do indeed have a shot of the fence. But I have to enhance the image so much that it becomes extremely noisy and grainy.

I can do a little better by making 10 or 12 copies of the image and compositing them all in Addition mode. But it's still very noisy.


But what if I take 12 *different* images of the fence, and then stack those?

Stacking different images allows you to increase the light level a lot with hardly any increase in the noise. You can even duplicate some layers: Figure 10-33 shows a stack of 24 layers total, two copies each of 12 different layers, all in Addition mode.

This technique reduces noise so well that you can boost the brightness even more using Levels or Curves without hurting the image much.

So grab a tripod (even one of the little plastic ones with legs a few inches long is fine if your camera is small enough), wander off to a dark place, and try taking some shots!





**Figure 10-33.** A 24-layer stack of the same fence is much less noisy. You can even see some stars!

## Increasing Contrast by Multiplicative Stacking

Those great shots that you see from NASA spacecraft aren't just snapshots downloaded straight from the camera. They take a lot of processing before NASA puts them up on public websites. (If you watch NASA TV, sometimes you can even see GIMP on screens at Mission Control! Of course, NASA uses a wide assortment of tools, not just GIMP.)

Figure 10-34 shows an original photo of Saturn's largest moon Titan, taken from the Cassini spacecraft. (You can find unprocessed images by looking for "raw images" on the Cassini website.) You can see that there's a little detail there, but it's hard to see it well. The contrast is very low.

But Cassini actually shot four of these Titan shots, and you can download all four and stack them.

Register the frames, and then set each layer's mode to Multiply (Figure 10-35).




Figure 10-34. Saturn's moon Titan, from NASA's Cassini spacecraft




Figure 10-35. Stacking several layers in Multiply mode boosts contrast, so you can see Titan's cloud details.

## Increasing Resolution by Averaging

When the Titan images were stacked, they didn't get sharper, merely more contrasty. But stacking, carefully done, can also increase sharpness. Technically, this is known as increasing *resolution*, the size of the smallest detail that can be distinguished in the photograph.

In recent years, video astrophotography has become very popular. Astrovideographers take cheap webcams or security cameras, modify them to mount on a telescope, and then take videos of bright objects such as the moon, Jupiter, or Mars. Each individual image is awful, and it doesn't make a very interesting movie. But when the video is split into its component frames—several hundred separate images—they can be stacked to make a very sharp photograph. Using these techniques, amateurs are now getting images as sharp as the professional observatory photos from a few decades ago.

First, the bad frames are removed. Turbulence in the atmosphere (even on a clear night) means that some images in these “movies” will be too blurry to be worth keeping. Then, the remaining frames are stacked using an averaging technique. Each layer is in normal mode, and the stacking is done with GIMP's Opacity slider. The bottom layer, of course, is at 100% opacity. Select the layer above it, and change the opacity to 50%. Now you're seeing an image that comes half from the first layer and half from the second: they're averaged. The resulting image won't be any brighter or darker than the components, but it will share pixels from each one.

Now make the third layer visible and select it. Set its opacity to 33%. That means that it'll be contributing only a third to the final image; the combination of the first two images will contribute two-thirds. So you're still seeing an average of the three.

Turn on the fourth layer, and set its opacity to 25%. And so on through all the layers. The rule you're following is that layer N gets opacity  $1/N$ :  $1/1$ ,  $1/2$ ,  $1/3$ ,  $1/4$ , and so on.

Of course, you wouldn't want to do this by hand if you had 200 layers. Astrovideographers usually use dedicated programs that handle both stacking and registration. Within GIMP, for a large number of layers, you'd be better off writing a script (you'll learn how to write simple scripts in Chapter 11).

## Stitching Panoramas

Ever find a beautiful sweeping view and feel frustrated that your camera's lens isn't wide enough to fit it all in?

Even if you have an SLR with replaceable lenses, sometimes you'll want to shoot vistas that are just too wide for the camera to capture in one shot.

Back in the film days (remember film?), I sometimes used to shoot several overlapping pictures, and then take the prints and tape them together and pin them up outside my office. They might have ended up looking something like Figure 10-36.

That might have been good enough a decade or so ago, but you can do a lot better than that now!



**Figure 10-36.** An old-style panorama, before GIMP

Combining several images into a panorama is called *stitching*. The process is straightforward, but it includes several steps, and requires some patience to get a result that looks truly seamless.

Let's take it one step at a time.

## Shooting the Images

Perhaps you already have a collection of images you'd like to stitch together. If not, here are some handy tips for taking a panoramic series. Some people find this stage the most intimidating part of the process. It's really not a problem if you follow a few guidelines.

### Overlap Photos for Panoramas

Overlap quite a bit. At first, you might be tempted to put the left edge of image 2 right where the right edge of image 1 was. There are two problems with that.

First, camera lenses (especially wide-angle lenses) distort the image (a phenomenon known as *lens warp*). The distortion is especially pronounced at the edges. If you use a lot of overlap, you'll be using mostly the centers of each frame where there isn't so much distortion.

Second, a lot of overlap means that if one image turns out to be bad, you can discard it and still make a usable panorama from the remaining images.

Overlap at least a third of each image (as in Figure 10-36), but more is better, especially if you're using a digital camera: "bits are cheap."

### Keep the Horizon Level

Keep the camera level: don't tilt it between shots. If you have a tripod, or a flat surface like a picnic table, that makes it easy.

If you have to hold the camera by hand (and let's face it, most of the time when you see a spectacular vista worthy of a panoramic shot, you probably won't have your tripod with you), try standing in a stable position and twisting your body as you take each shot.

Pay attention to the horizon in the camera's viewfinder: try to keep it parallel to the top and bottom of the frame. If you accidentally tilt the camera, you can use GIMP's Rotate tool to make it level after the fact. However, rotating images so that they all match well enough to make

a panorama is more difficult than it sounds, and you can save yourself a lot of work by shooting them level to begin with.

It's less important that you keep the horizon at the same height for each photo. You can have more sky in one shot, more ground in another, and the panorama will still look fine. The downside: to get a rectangular picture you may have to trim more from the top and bottom than is ideal.

## Match Exposure Levels

If you can shoot all the images at the same exposure setting, that may save a little work later. Otherwise, the individual photos may vary in color or brightness. Notice the different colors of the sky and the ground between the left and middle pictures of Figure 10-36.

Some cameras offer a panorama setting which does this automatically, or manual exposure controls. If yours doesn't, you may be able to simulate it by *metering* on a point somewhere around the middle of your panorama: hold the shutter button halfway down, rotate the camera (still holding the button halfway), and press the button the rest of the way to take the picture. Repeat as necessary.

This isn't crucial. You can adjust brightness in the GIMP, of course. And a panorama whose components have slightly different brightness levels may still look okay. Try it both ways and see whether it's worthwhile for you.

Once you've shot all your images and uploaded them to GIMP, it's time to start stitching.

## Decide on a Resolution

Although it's normally best to work at the full resolution of your camera and scale the final image down after you're finished, that's not necessarily true for a panorama.

Panoramic images combine a lot of photos into one humongous GIMP image. Using three-megapixel or larger images for a big panorama (say, five or more images) can take up so much memory that it slows the machine to a crawl. You might use full resolution if you need an image that large (for instance, if you plan on printing to a banner three-feet long). But otherwise, if you notice slowness, consider scaling down a bit.

Once you know the resolution of the individual images, how many of them there are, and the overlap you used, you can calculate how big an image you'll need to hold the panorama.

## Calculate Your Expected New Image Size

The next step in building a panorama is to make a new image of the right size. What is the right size? That's something you'll have to figure out.

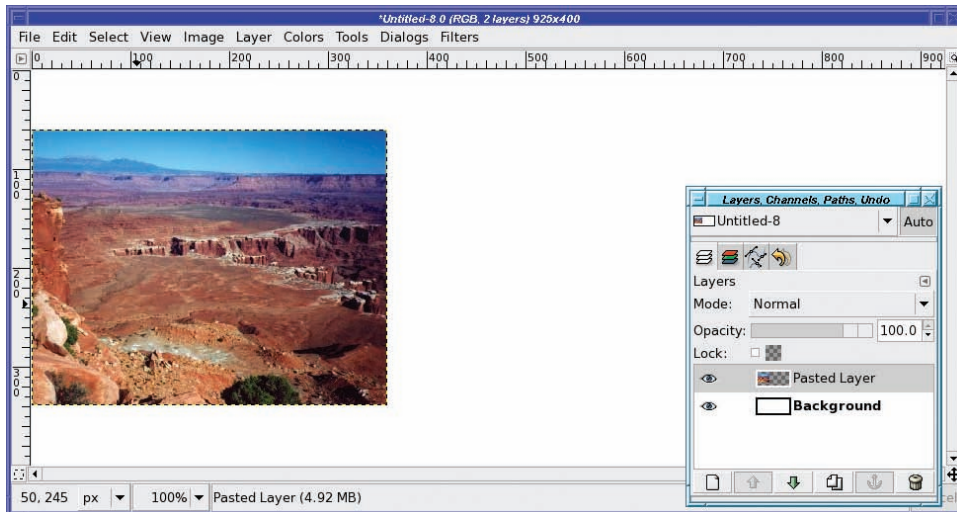
The easy way out is to simply take the width of your images and multiply by the total number you'll stitch together. There's not much penalty for starting with a new image that's too big—you'll just crop it in the end anyway. (GIMP may complain if you try to create a new image larger than the parameters set in your preferences. The default, which you probably didn't change, is 64MB. Should you need to change this setting, you'll find it in Preferences under *Environment*.)

If you're feeling lucky, subtract whatever you think you can get away with, say 20 to 25%. Don't stress over it. If it's too big, crop it later. If it turns out you need more space, a simple *Image* ► *Canvas Size* will fix it in a jiffy.

Once you've decided on a size, create the new image. I'll refer to this image as the "panorama image."

## Load the First Two Images

Open the first image as a new layer and move it (using the Move tool) all the way to the left side of the image (assuming your panorama is left to right). The panorama image will look something like Figure 10-37.



**Figure 10-37.** The first image is loaded.

Next, load the second image as a new layer, just as you did the first, and use the Move tool to move the new layer to where it approximately overlaps the first one. Don't worry about getting it perfect . . . yet.

## Make a Gradient Layer Mask

How do you get rid of the sharp edge between the first picture and the second? By using a layer mask.

In the Layers dialog, right-click on the second layer and choose *Add Layer Mask...* Initialize it to *White (full opacity)*, the default.


Now use the Gradient tool to draw a gradient on the layer mask. It should be white everywhere to the right of the overlap, shading to black at the left edge of the second image. If you have the normal black foreground and white background, and the Gradient tool is in normal mode (not *Reverse*), this means starting your drag from the left edge of the second image, and then dragging right for some distance, not all the way to the point where the first image ends. You'll end up with a result like Figure 10-38.

---

**Tip** It helps to use the Control key while dragging your gradient, to help get a gradient that's exactly horizontal.

---





**Figure 10-38.** Draw a gradient on the layer mask so that the left edge of the second image fades into the first.

## Fine-Tune the Position of the Second Image

Now that the second image has transparent areas, you can position it more accurately. Switch to the Move tool, but there's a trick: if you simply click and drag with the Move tool, you'll move the layer mask rather than the layer itself, since the mask is still selected in the Layers dialog (indicated by the white boundary around it in Figure 10-38). Click on the layer preview in the Layers dialog to select the layer and not just its mask.

Zoom way in so that you can see details, and scroll so you can see the region of the image where the two layers overlap. Then use the Move tool to move the second image around until it matches the first image as well as possible. (Don't forget about the arrow keys.)

It'll usually be obvious when the images mesh. When you have two overlapping images (one partially transparent) that don't quite match, the image will look a bit blurry, as though you had a few too many beers before beginning the stitching process. When you find the right position, suddenly the image will seem to snap to focus, and all the blurriness and double vision will go away.

The catch? It won't happen everywhere in the image. Most of the time, the two images won't match up perfectly. If you line them up in the center, the top and bottom won't match. You're seeing lens warp. If you overlapped the images a lot (50% or more) or zoomed in when you shot the images, you'll see less lens warp, but there will always be some.

Usually, you're best off matching two layers at the vertical midpoint, and not worrying about the top and bottom (you can fix that in the next step). But if your panorama has a pronounced horizontal line somewhere—say, the horizon in an ocean shot—you may want to use that as the location to match.

## Adjust the Layer Mask

The gradient in your layer mask was great for a first step, but now it's time to eliminate parts of the image that don't match well. You don't want to replace the gradient, just add to it.

In this case, the first priority is getting rid of those mesa tops that don't line up. To do that, a bit of black added with the Paintbrush tool in the upper-left of the second image's layer mask will make the ghostly floating mesa disappear. (Remember, black in the layer mask makes that part of the image invisible.)

Watch the double images gradually go away as you scribble—it's fun.

If you need to get rid of a whole corner of an image, you can add another gradient to the existing gradient. Use the Mode menu in the Gradient tool: any of the darkening modes, such as Multiply, Subtract, or Darken Only, will do (they'll give slightly different results at the edges where the two gradients meet, but you probably won't see the difference).

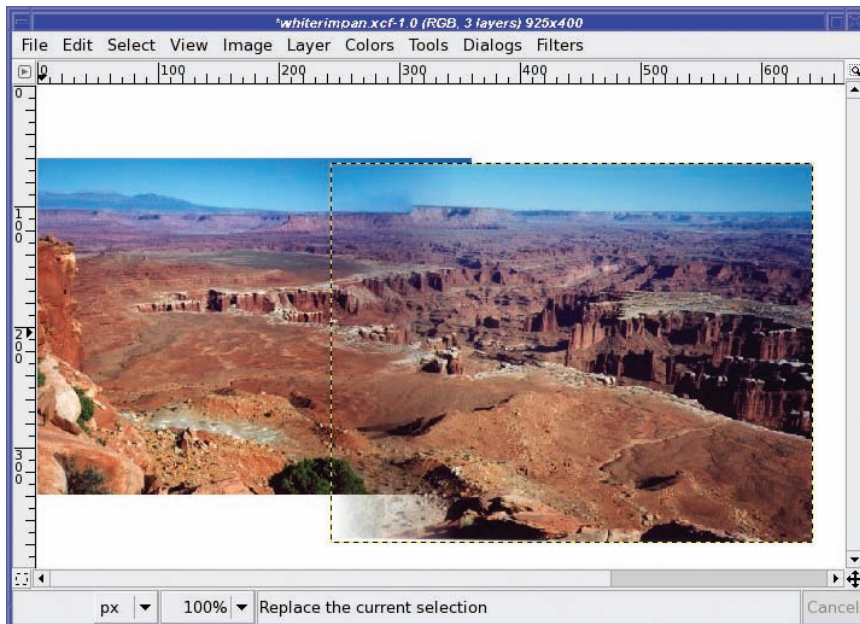
---

**Tip** Remember that you can use *Show Layer Mask* from the Layers dialog (or Alt-click or Alt-Shift-click on the layer preview).

---

When you're finished adjusting the mask, click on the layer preview in the Layers dialog so the mask won't remain selected. That way, you won't be surprised if you try to do anything with the layer later.

Figure 10-39 shows the mesa-top area cleaned up (I cheated a little and made a new butte). Notice that the sky now looks patchy. Don't worry, you can clean up the sky later.



**Figure 10-39.** The area around the mesa top is cleaned up. The sky still needs work.

## Adjust with Other Tools if Necessary

Most of the time, you can get what you need (except the sky) by drawing on the layer mask. But in a few cases, you'll have something—perhaps a branch, or a road—that just doesn't meet in the two images.

Sometimes, you can use other tools to fix this. Of course, you can draw directly on the image, using your cloning and smudging skills from Chapter 6. But first, try IWarp (described in Chapter 7, “Distorts”). Sometimes careful use of IWarp can do exactly what you need.

## Add the Rest of the Images

You're done adding the second image! Now, repeat the steps for each of the pieces of your panorama: load the image, add a layer mask, move it into place, adjust the layer mask, and make any additional adjustments you need.

## Final Adjustments

Ironically, skies are particularly hard to get right in panoramas. Differences in exposure seem much more obvious against a clear blue sky. So go ahead and cheat: select the sky using your favorite selection technique, and then make a sky from a gradient between two shades of blue, just like in the “Decomposing to HSV” exercise in Chapter 8.

Finally, crop the panorama. The top and bottom edges of the component images probably won't quite line up. It's your choice: you can leave the edges ragged (Figure 10-40), to show off its panoramic nature, or you can crop to an even rectangle so it looks like the whole thing was taken in one shot.



**Figure 10-40.** *Crop the panorama. Leave the ragged edges if you like: sometimes it's fun to show how many images went into your creation.*

## Summary

Now you know all sorts of ways to composite several images. No book could possibly list all the amazing things you can do compositing images in the GIMP, but now you know the important tools and a collection of different techniques. You should have a good solid base from which to begin experimenting with your own images.

You'll also be able to apply techniques you find in web tutorials. If you experiment with these ideas, and try them on the sorts of images you use yourself, you'll be an expert in no time.

But for now, take a break from fiddling with images by hand. During the course of this chapter, you've done quite a lot of repetitive operations which would have been so much easier if they were automated. In the next chapter, you'll learn how to automate GIMP processes yourself, using plug-ins and scripting.



# Plug-Ins and Scripting

**T**he GIMP is a supremely flexible image-editing program. But sometimes you may need to do something unusual, something that isn't built-in. Or maybe you just need to do the same repetitive operation over and over on many different images.

Either way, don't give up just because the function you need isn't in the menus anywhere. GIMP's flexible architecture means it's very easy to add plug-ins and scripts to do nearly anything you need to do. There are dozens, maybe hundreds, of plug-ins readily available. But if you can't find one to do what you need, you can even write your own, using a choice of several different languages.

Many plug-ins are written in the C programming language, the same language used to write GIMP itself. But other plug-ins are actually scripts.

*Scripts* are a special type of plug-in, a type that doesn't need any special *compiler* software. The GIMP supports three scripting languages, *Script-Fu*, *Python*, and *Perl*, in varying degrees. You'll learn more about the differences later in this chapter.

Scripts are especially easy to write, which has three benefits:

- It's easy for other people to write them, so there's a good chance you can find a script to do what you need.
- It's easy to read them, so if you can find a script that's similar to what you need, you might be able to modify it for your purposes.
- It's not very hard to write a script from scratch.

So let's take a look at what all this means. You'll learn how to find and install plug-ins, and how to find out which plug-ins you already have installed. You'll also learn how to find scripts, install them, read them, and modify them. And finally, you'll learn how to write a new script from scratch to automate some simple but useful operations. The chapter will cover the following topics:

- Installing plug-ins
- Writing GIMP scripts
- User interface options for scripts
- A Python script

- Finding out how to do things: The GIMP PDB
- A Perl script
- Batch processing of multiple files
- Writing a C plug-in

## Plug-Ins

Strictly speaking, anything that isn't built into the GIMP's core libraries is a plug-in. You've probably already used lots of GIMP plug-ins: nearly everything in the *Filters* menu is a plug-in. About half of the *Colors* menu and most of the *Xtns* menu in the Toolbox are also plug-ins.

### The Plug-In Browser

How do you know which plug-ins you already have installed? With the Plug-In Browser (Figure 11-1), accessed through the Toolbox menu: *Xtns* ► *Plug-In Browser...*

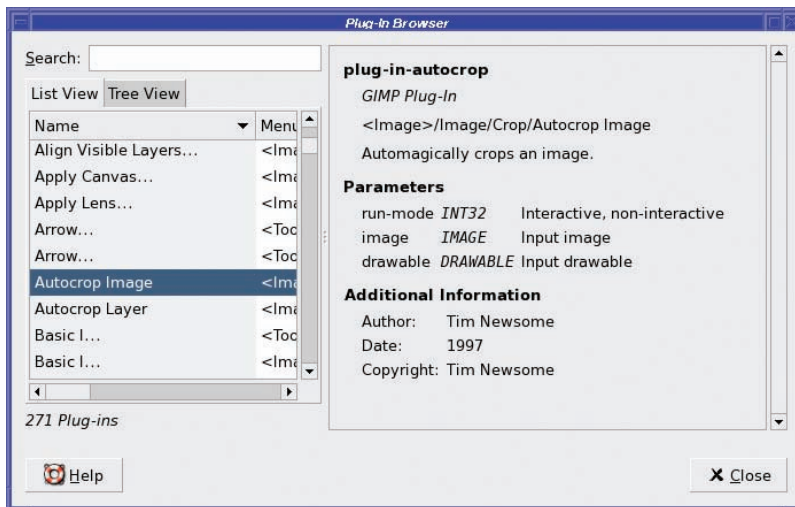



Figure 11-1. The Plug-In Browser

Yikes—there are a lot of plug-ins! Bet you didn't know you had that many installed.

Fortunately, there's a *Search* field. Type in a search term, such as **crop** (or even part of a term, like **rop**), and GIMP will show you only plug-ins that contain that term.

The Plug-In Browser is a good way of finding functionality without searching through all the menus. If you're pretty sure that there's some sort of automatic crop function, click on the *Tree View* tab and enter **crop** in the *Search* field. All plug-ins that include the name “crop” will come up in the menu, and you can see where they are in the menu structure, as in Figure 11-2.





**Figure 11-2.** The Tree View shows GIMP’s menu structure; entering a term in the Search field limits the number of plug-ins shown.

The *Image Type* field tells you which kinds of images the plug-in accepts. Some plug-ins will only work on RGB images, and will be grayed out if the current image is indexed or grayscale. Other plug-ins may work only on indexed, only on grayscale, or only on RGBA (RGB images with an alpha channel). A “\*” is a wildcard: “RGB\*” means that both RGB and RGBA images will work, and “\*” by itself means the plug-in will accept any type of image.

The right half of the dialog gives you more information about the plug-in, including its location in the menus, a description of the plug-in, the parameters it takes (these are usually the same as the values you’ll be able to change in the plug-in’s dialog) in case you ever wanted to use it in another plug-in, and information about who wrote the plug-in.

## Finding Plug-Ins

What if you’ve searched the Plug-In Browser and there’s nothing there? Might someone else have developed a plug-in to do what you need?

You have two options to find such a plug-in: the GIMP’s Plug-In Registry, and a web search.

### The Plug-In Registry

GIMP’s official Plug-In Registry is located at <http://registry.gimp.org>. Anyone who develops a plug-in or script for GIMP can register it there. As a user, you can search by language, by GIMP version, or by categories such as *Distorts*, *File*, *Artistic*, and so on. If you aren’t sure which category would cover the function you’re looking for, just search by name and hope you find something.

The Plug-In Registry is theoretically a good resource. But it tends not to be very well maintained: the GIMP versions it lists tend to run behind, and searching doesn’t always work well.

Try the Plug-In Registry when you’re looking for a way to do something in GIMP. But if you don’t find anything there, don’t despair: try a web search.

## Finding Plug-Ins Using a Web Search

A web search often works much better for finding GIMP plug-ins. Lots of people develop plug-ins and never get around to listing them on the registry. Maybe they don't know about the registry, or they don't think their plug-in is quite polished enough to list there, even though it might be good enough for your purposes.

Use “gimp” as one of your search terms, along with the functionality you're looking for, such as “gimp redeye.” If you get too many hits (sometimes you'll get a lot of mailing list traffic discussing how to do things by hand), you could try adding “plug-in” or “plugin” or “script” to your search query, for example, “gimp redeye plugin OR script.”

## Installing a Plug-In

With any luck, the plug-in you find will come with installation instructions. But some don't. Here's what you need to know about installing plug-ins.

You first need to know what type it is: whether it's a C plug-in or a script. Usually the page where you found the plug-in will tell you. In addition, a script will usually have an extension such as .scm (Script-Fu), .py (Python), or .pl (Perl), while a C file will end in .c.

If the file you downloaded has the extension .zip, .tar.gz, or .tar.bz2, that means it's part of a compressed archive. You must unpack the archive first (using a program such as *unzip* or *tar*) before you can install the plug-in.

## Installing a Script

Installing a script is easy. For a Script-Fu, just take the *file.scm* and copy it to the folder named “scripts” inside your GIMP profile. For Python or Perl, copy the file to the “plug-ins” folder instead of “scripts.” If you don't remember where your GIMP profile folder is, check the *Folders* category in the *File* ➤ *Preferences* window.

That's all! The next time you start GIMP, the new script should register. If GIMP is already running, you can tell it to look for new Script-Fu scripts from the Toolbox menu: *Xtns* ➤ *Script-Fu* ➤ *Refresh Scripts...* If the script is Python or Perl, you'll have to restart GIMP.

---

**Tip** If you have the *gimptool* program installed, you can use it to install Script-Fu scripts. In a terminal window, type **`gimptool --install-script file.scm`** (that's two dashes together). For Python or Perl plug-ins, use **`gimptool --install-bin file`**.

This shortcuts the need to find your GIMP profile, if you don't mind typing commands. You'll need *gimptool* for installing C plug-ins anyway (see below).

---

## Installing a C Plug-In

Installing a C plug-in is no more complicated—if the plug-in has already been “compiled” for your platform. C is a *compiled language*, which means that in order to run anything written in C, it has to be translated (compiled) into machine language first. The compiled version is referred to as an *executable* or a *binary*.

Compiling makes a C plug-in run much faster than a script, since the compiler speaks the machine's native language. However, it also means there's an extra step in creating something that can run, and it has to be done separately for each platform.

If there's a binary available for your operating system, just take that file and move it to the plug-ins folder inside your GIMP profile. Alternately, if you have `gimptool` installed, you can use `gimptool --install-bin`. The next time you restart GIMP, the new plug-in should show up in the menus.

---

**Tip** Where did the new plug-in show up? A lot of plug-in developers forget to tell you where to look in the menus to find their plug-in once it's installed. Since GIMP has so many features already, searching for something new could take quite a while! Fortunately, you can use the Plug-In Browser to figure out where it appears.

---

However, since many plug-in developers don't have C compilers for all three platforms, often they simply make the C source code available, and you have to build (compile) it yourself. You'll need some supporting programs in order to do that.

## Installing a Gimp C Development Environment

On Mac and Linux systems, installing a C compiler is easy. The GNU C compiler (`gcc`) comes on the Mac OS X Development CD, and is included with most Linux distributions. If it's not installed, it can usually be installed easily using the software update tools provided with your system. You'll probably need a few additional packages, in particular `gimptool`. It's sometimes included with GIMP, but if not, look for a package called something like `gimp-dev` or `gimp-devel`.

You'll also need development packages for X11 and `gtk`. Most packaging systems should automatically install these when you install the GIMP development package, but if you get errors, these packages may be what you're lacking.

On Windows it's a bit harder, since Windows doesn't come with a C compiler. If you don't already have a compiler installed, you're probably best off with MinGW, the Minimalist GNU for Windows packages. Find them at <http://www.mingw.org/>. You will need both MinGW and MSYS.

In addition, even if you already had a compiler, you'll need the development libraries for GTK+, Glib, and other related features. You can get them by clicking on "GIMP for Windows" on the main [gimp.org](http://gimp.org) site. Look for packages with "-dev" in the name.

## Building a C Plug-In

Once you have the compiler and libraries installed, it's time to build your plug-in. You'll need a terminal window open so you can type commands: the *Command Prompt* program should work on Windows systems, *Terminal* on Macs, or any terminal client on Linux. Then type:

```
gimptool --build filename.c
```

where `filename.c` is the name of the plug-in file. Remember: two dashes, not just one.

If this gives errors, it may be that you're missing libraries. Double-check that you have all the requirements installed.

Once the compile completes without errors, it's time to install it. Type:

```
gimptool --install filename
```

where filename is the file you built in the previous step. (Some very complex plug-ins may need to install more than one file, and use their own installers, usually run from a “Makefile,” so the user need only type **make install**. Creation of Makefiles and installers is outside the scope of this book.)

## Writing GIMP Scripts

Now it's time to take a closer look at GIMP scripts.

The GIMP can be scripted in three different languages: Script-Fu, Python, and Perl. They each have their advantages and disadvantages.

### Script-Fu (= Scheme = Lisp)

Script-Fu (a pun on “kung-fu”) is the GIMP's built-in scripting language. Script-Fu has one big advantage: it is always available in every version of the GIMP. If you want to write a script to distribute widely, Script-Fu is the best choice. You can be fairly sure everyone will be able to use it.

In addition, most scripts distributed with the GIMP are written in Script-Fu. That means that you can find examples already installed on your system for all sorts of useful techniques. Since the easiest way to learn scripting is to copy an existing script (as you'll see in a moment), that's a big advantage.

The downside of Script-Fu is that its syntax is somewhat difficult for most people. Script-Fu is actually just the GIMP's term for a language called Scheme, a dialect of Lisp. (To be even more specific, it's a dialect of Scheme called SIOD, and you can find tutorials and documentation on its syntax with a web search for “SIOD.”)

*Lisp* is a language invented in 1960 for mathematics and artificial intelligence research. The name stands for “*list processing*,” because it's very good at processing lists of items. But Lisp programmers joke that it actually stands for “Lots of Irritating Stupid Parentheses,” a phrase which you will readily understand after looking at some Script-Fu examples.

Don't be too scared off by the parentheses. Often, you can copy bits and pieces from existing programs to make simple scripts without needing to understand much about Lisp. But if you really dislike the syntax, you may be better off with one of the GIMP's other scripting languages.

### Python

Python is a much more modern language, created in 1990. It has a clear, straightforward syntax and is a very friendly introduction to general-purpose programming. It's also well integrated into the GIMP, so it's fairly easy to write gimp-python scripts.

The only downside is that Python didn't become part of the standard GIMP install until quite recently, and was a bit tricky to install on Windows. That's been fixed now—gimp-python comes as part of GIMP 2.4—though you do need the Python language installed on your system: you can get it from <http://python.org> if you don't already have it. You may also need the pygtk package, from <http://pygtk.org>.

You can easily tell whether you have gimp-python enabled: check for a Python menu in the Toolbox under *Xtns*.

If you're just learning scripting and you have gimp-python installed, it's well worth trying it out.

## Perl

Perl is a well established and popular scripting language, commonly used on web servers and for system administration.

Perl used to come standard with GIMP, at least on UNIX systems. But somewhere along the line, it became less popular for GIMP scripting, and today it's not included as part of the standard GIMP package. You can see whether you have gimp-perl installed by looking for a *Perl* or *Perl-Fu* menu under *Xtns*.

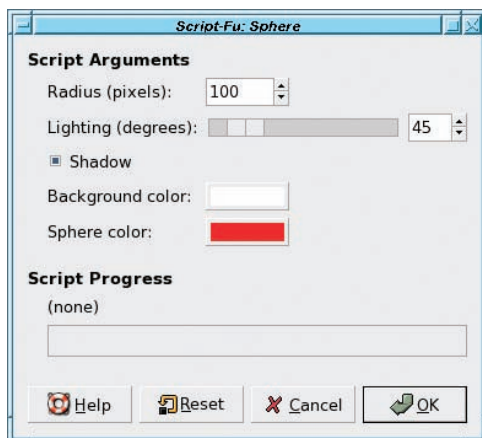
It's still possible to install gimp-perl as a separate package (look for it at [gimp.org](http://gimp.org)), and there are still plenty of gimp-perl scripts in the Plug-In Registry. If you're familiar with the Perl language and prefer it to Python or Scheme, then there's no reason not to use Perl as your introduction to GIMP scripting. If you're not already a Perl fan, though, you're probably better off sticking to Python or Script-Fu.

## Reading and Modifying a Script


You don't have to be a programmer to write simple GIMP scripts, or to modify existing scripts, though of course programming experience will give you a head start.

It's easiest to begin with an existing script and modify it. A quick run through the process will show you how it's done.

Let's dig deeper into something that always frustrates me. The GIMP has a useful little Sphere script, in the Toolbox *Xtns* menu under *Misc* (in earlier versions, look for *Misc* inside *Script-Fu*). It shows you the dialog in Figure 11-3, and produces a nice sphere like the one in Figure 11-4.



**Figure 11-3.** *The Sphere script's dialog*



**Figure 11-4.** *The sphere generated by the built-in Sphere script*

You might want to take one of these spheres and paste them into a drawing—wouldn't it make a nice Christmas tree ornament or poodle nose? But you can't! The sphere generated by the script always has a solid background.

If you want to paste the sphere into another image, you either have to use your selection skills to get rid of the background, or make your own sphere by hand (as you learned to do in Chapter 9). That's a lot of work. But the GIMP is an open source program, and that means you aren't limited to the choices the developers made: you can make it do anything you want. Why not just make it do the right thing?

### Find the Original Script

The first step is to locate the existing script. You can find where GIMP scripts are stored by looking in the Preferences window (*File* ► *Preferences* from the Toolbox): expand the *Folders* category and click on *Scripts*.

The GIMP should show you at least two folders. One is the place where you can put personal copies of scripts; the other is in a system directory where GIMP has its own files installed.

Make a note of these two locations. You'll use them both.

Now find the script itself (in the system script directory). Usually it will have a .scm (short for “scheme”) extension. You generally have to guess what the file name might be, but in this case it's not too hard: there's a file there called *sphere.scm*.

### Edit a Copy of the Script

It's not a good idea to edit GIMP's existing scripts in place—what if something goes wrong?—so be safe and copy *sphere.scm* to your personal scripts folder.



---

**Caution** There's a trick to copying the script on a Mac. Since it's inside an "Application," double-clicking the folder will launch Gimp.app! Instead, either right-click on the application icon, or use Control-click if you have a one-button mouse. You'll be presented with a pop-up dialog, where you can select *Show Package Contents*. This will open a normal folder from which you can navigate to your scripts folder in the usual fashion. Alternately, for the adventurous, you can use the Unix copy command `cp` from Terminal.app.

---

Give it a new name, like `mysphere.scm`, and open it with a text editor. You need a program that handles plain text, not a word processor.

### WHY A TEXT EDITOR?

Most word processors save to a format designed to preserve text fonts, styles, and layout information. The result may not be readable by *anything* other than the original word-processing program, and GIMP probably won't understand it at all.

It's better to just use something that naturally reads and writes plain text. A simple program like Windows WordPad is fine if you don't already have a favorite text editor. (Why not Notepad? It does handle text, but only the sort written in Windows. In particular, it's not very good at handling the various kinds of line breaks—"returns"—used on other computer systems.)

On Mac, TextEdit works perfectly. Normally, TextEdit is smart enough to open a `.scm` file as plain text, and save it the same way. But you should check anyway to make sure—if you can see *Make Plain Text* in the *Format* menu, use it!

On Linux, you have quite a variety of text editors available—such as `joe`, `vim`, `emacs`, `ee`, `gedit`, `keddit`—and you probably have a favorite already.

If you find yourself writing a lot of plug-ins, a good programmer's editor can make it easier by highlighting different parts of the program in different colors, handling indentation for you, matching parentheses, and other niceties.

Take a deep breath, open `sphere.scm` in your text editor, and follow along on a tour through the script.

### Make the Script Your Own

All GIMP scripts use a *registration procedure* to tell GIMP the name of the script and where to put it in the menus. The name of the file doesn't figure in.

Since you copied the script to a new place but haven't changed the name it uses when it registers, GIMP now has two scripts called *Sphere*, in exactly the same place in the menus. That's no good—this means that when you call one of them, you don't know which of the two scripts you'll get. To fix that, you need to change the name of the script—not just the file name, but the names it registers inside the file.

Skip to the end of the file and look for `script-fu-menu-register`:

```
(script-fu-register "script-fu-sphere"
  "_ _Sphere..."
  "Simple sphere with a drop shadow"
  "Spencer Kimball"
  "Spencer Kimball"
  "1996"
  ""
  SF-ADJUSTMENT _ "Radius (pixels)"      '(100 5 500 1 10 0 1)
  SF-ADJUSTMENT _ "Lighting (degrees)"    '(45 0 360 1 10 0 0)
  SF-TOGGLE      _ "Shadow"              TRUE
  SF-COLOR       _ "Background color"    '(255 255 255)
  SF-COLOR       _ "Sphere color"       '(255 0 0))

(script-fu-menu-register "script-fu-sphere"
  _ "<Toolbox>/Xtns/Misc")
```

`script-fu-register` has several jobs. First, it sets the name of the routine that does the work. In this case, that's `script-fu-sphere`. You'll see how that's used in a moment.

Right after that is the name which GIMP will show the user: in this case, `_ _Sphere...`. That's the name that shows up in the menus when you call the function. It's a good idea to change that now, so you can tell the difference between your version of the script and the one that's built-in.

What are those underscores? They're both optional; you may include them or not, as you choose.

The first one, outside the quotation marks, helps with translating the script into other languages. Someone translating the Sphere script will make a "catalog" of all the strings in the script and their translations. The underscore is an indication that this is a string that might need to get translated. Behind the scenes, it also helps implement the actual display of the translated strings so the user will see the right language. You can leave it off, but then it will be harder for speakers of other languages to adapt your script.

The second underscore has to do with menus: it's a way of accessing the function by using the keyboard. (If you have the *Misc* menu showing, and there's an underscore before the S in *Sphere*, then pressing S will run that function.) As a beginner script writer, you're fine leaving this underscore off (though it usually won't hurt anything if you leave it there).

So change `script-fu-sphere` to something like `script-fu-my-sphere`, and change `_ _Sphere...` to `_ "My Sphere..."` (notice that the underscore outside the quotes has been left in, so the script can be translated, but the one inside has been removed). Be sure to change the `script-fu-sphere` in `script-fu-menu-register` as well. So now they look like

```
(script-fu-register "script-fu-my-sphere"
  _ "My Sphere..."
  "Simple sphere with a drop shadow"
```

and

```
(script-fu-menu-register "script-fu-my-sphere"
  "<Toolbox>/Xtns/Misc")
```

Since you've changed `script-fu-sphere` to `script-fu-my-sphere`, you also need to change it where that function is defined: in this case, at the beginning of the script.

```
(define (script-fu-my-sphere radius
        light
        shadow
        bg-color
        sphere-color)
```

You may want to try *Xtns* ► *Script-Fu* ► *Refresh Scripts* at this point, to see if your new “My Sphere” script appears in the *Xtns* ► *Misc* menu of the Toolbox.

## Try Making a Change

To make minor changes, you don’t need to know much about Script-Fu. You can usually guess where to start. In fact, learning *how* to guess is one of the more important skills!

In the sphere.scm case, the goal is to remove that background color. A reasonable first guess would be to start at the beginning and search for the word “background.”

The first place you find that is:

```
; bg-color: background color
```

Anything that begins with a semicolon in Lisp is a *comment*. It’s just a note the programmer adds to explain aspects of the script. A comment can start in the middle of a line, too: anything from the semicolon to the end of the line is part of the comment. Adding comments is always a good idea when you do anything tricky that you might not remember six months from now, or which you think you might need to explain to someone else.

This comment is hinting that `bg-color` is associated with background color somehow. That’s your clue that in addition to looking for “background,” you should probably also be on the lookout for “bg-color” as you continue working through the script.

The next occurrence of either of those strings comes in this block:

```
(gimp-context-set-foreground sphere-color)
(gimp-context-set-background bg-color)
(gimp-edit-fill drawable BACKGROUND-FILL)
```

Without knowing Script-Fu, you can probably take a guess at what this is doing. It’s setting the foreground color to `sphere-color` (the color you set in the Sphere dialog) and the background color to `bg-color` (also set in the dialog). Then, it’s filling the “drawable” (whatever that is) with the background color, using something similar to the *Fill with BG color* item in GIMP’s *Edit* menu.

What’s a *drawable*? In a GIMP plug-in, a drawable is anything you can draw on. Often, it’s a layer, but it could also be a layer mask or channel. In the Sphere script, if you search backwards for “drawable” (you’ll have figured out by now that knowing how to search in your text editor is an important part of scripting!), you find this:

```
(drawable (car (gimp-layer-new img width height RGB-IMAGE
                "Sphere Layer" 100 NORMAL-MODE)))
```

A long and somewhat complicated line! But if you scan over it, you’ll notice a good clue: `gimp-layer-new` suggests that this is creating a new layer. In fact, “drawable” is the first and only layer used in the Sphere script. (What “car” means will be explained in a bit.)

So that `gimp-edit-fill` line is the first suspect. Try commenting it out by putting a semicolon in front of it:

```
;(gimp-edit-fill drawable BACKGROUND-FILL)
```

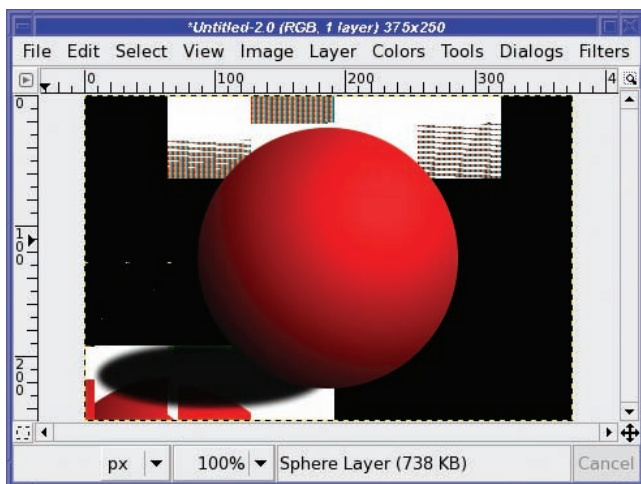
## Try Your Fix

Now it's time to see if your fix worked. But first, you need to tell GIMP to reload scripts, so it will see your changes. (*Xtns* ► *Script-Fu* ► *Refresh Scripts*, as you learned earlier.)

If you changed the name of your script, GIMP should now find *My Sphere...* (or whatever you named it) in the *Misc* menu right next to *Sphere...* Run it.

The dialog pops up. There's still a chooser for *Background color*: there. That's okay: you didn't do anything to remove it. (You can worry about that later.) Go ahead and click OK.

Whoops! You probably see something like Figure 11-5. Obviously you have a little more work to do.



**Figure 11-5.** *Sphere taken over by cubists*

## What Went Wrong? Some GIMP Peculiarities

There are two problems. First, the background was supposed to be transparent, but it's not. Second, what's all that stuff filling the space behind the sphere?

First things first. When you expect transparency but don't see it, your first question should be whether the image might be lacking an alpha channel. Unfortunately, GIMP doesn't provide a straightforward way to check. Try right-clicking on the layer in the Layers dialog: if *Add Alpha Channel* is enabled, then the layer doesn't already have an alpha channel (which indeed is the case for the sphere layer).

---

**Tip** Current versions of GIMP 2.x also show layers without alpha in boldface type in the Layers dialog, so you may be able to use that as a quick check.

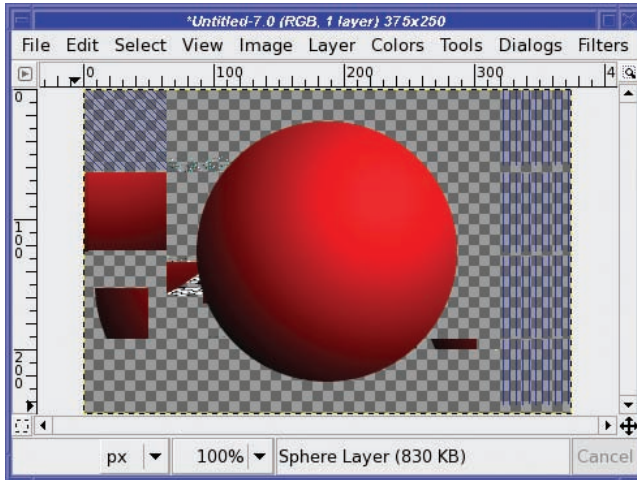
---

How do you change the script to include an alpha channel in the drawable layer? Change RGB to RGBA (the A, of course, stands for *Alpha*).

```
(drawable (car (gimp-layer-new img width height RGBA-IMAGE
                "Sphere Layer" 100 NORMAL-MODE)))
```

Save the file, *Refresh Scripts*, and then run *My Sphere...* again.

Progress! Now the background is transparent. Well, mostly . . . but there's still a lot of random garbage there, as in Figure 11-6, different every time. What's going on?




**Figure 11-6.** Now the background is transparent, except that it's full of random garbage.

One thing you need to know when writing GIMP scripts is that when GIMP creates new layers, it doesn't automatically clear them. GIMP doesn't know whether you're going to want a transparent layer or a filled layer, so it doesn't waste its time with one or the other. Consequently, the layer might end up being filled with whatever happened to be in that part of your computer's memory.

That sounds annoying, but all it means is that every time you create a layer, you should immediately either clear it, or fill it—do the equivalent of *Edit* ► *Clear* or *Edit* ► *Fill with ....* Since you commented out a call to `gimp-edit-fill` which filled with the background color, replace it now with a call to `gimp-edit-clear`:

```
(gimp-edit-clear drawable)
```

Another *Refresh Scripts* and you're ready to test it. Voilà! This time it works, and you see something like Figure 11-7.



**Figure 11-7.** Success! A sphere against a transparent background.

Hooray! You now have a script to make a sphere that you can paste anywhere.

### Bringing the Shadow Back

But wait! The original script had a shadow. You didn't remove the shadow from the script—yet it's gone. Why?

To find out, search for “shadow” in the script. You'll find a lot of operations involving `shadow-x`, `shadow-y`, `shadow-w`, and `shadow-h`: after you've read or written scripts for a while, seeing the combination `x`, `y`, `w`, and `h` will immediately make you think of *x* coordinate, *y* coordinate, *width*, and *height*. Read past those, and you get to the following:


```
(gimp-ellipse-select img shadow-x shadow-y shadow-w shadow-h
                           CHANNEL-OP-REPLACE TRUE TRUE 7.5)
(gimp-edit-bucket-fill drawable BG-BUCKET-FILL MULTIPLY-MODE
                        100 0 FALSE 0 0))
```

The script selects an elliptical area, and then fills it—but it fills in Multiply mode. Multiply mode of a dark shadow against a white background makes sense; but it won't show anything against a transparent background.

Change `MULTIPLY-MODE` to `NORMAL-MODE`, re-run the script, and you get Figure 11-8.

If you intend to continue using your sphere script, you'll probably want to make a few more changes. Since the background color isn't used any more, remove it from the dialog, both in the `script-fu-register` command and in the parameters to the `script-fu-sphere` routine—or perhaps change the script to create a new layer with the background color, so you can turn it on or off as you choose. You can keep tweaking and adding as many new features as you like.





**Figure 11-8.** *The shadow in normal, not Multiply, mode*

## Anatomy of Script-Fu: A Stroll Through the Sphere Script

Let's go through the original Sphere script now, and try to understand what it's doing. It uses a lot of techniques that will work when writing your own scripts.

Lisp expressions are called *lists*, and they're always delimited by parentheses. Items inside the list are separated by spaces. These parenthetical expressions can be “nested” as deeply as necessary.

A Lisp function is just another list. The function is the first element of the list, and any parameters the function takes make up the rest of the list. Where some computer languages might say,

```
gimp-image-new (width, height, RGB)
```

Lisp would say,

```
(gimp-image-new width height RGB)
```

Arithmetic operations and conditional expressions such as “if” and “while” are just ordinary functions in Lisp, so they're expressed as lists too. For example, to add 3 to the multiple of 2 and 5, you would say:

```
(+ 3 (* 2 5))
```

This is called *prefix notation*. Read expressions like this from the inside out. In this case, first 2 and 5 are multiplied, and then that result is added to 3.

If you have a text editor that lets you put the cursor on a parenthesis and find the one that matches it, that can be very helpful when working on Lisp programs.

## The script-fu-sphere Function

The Sphere script begins by defining the routine that does the work:

```
(define (script-fu-sphere radius
          light
          shadow
          bg-color
          sphere-color)
```

The values *radius*, *light*, *shadow*, *bg-color*, and *sphere-color* are called *parameters*. Nearly every script has parameters: they're values that control exactly what the script will do. In Script-Fu, they usually correspond to the items you can change in the script's dialog.

Often, they represent sizes (the radius of the sphere) or colors (of the sphere and the background). Sometimes they are true/false values: *shadow* is such a value, and corresponds to a checkbox in the Sphere dialog. *Light* is a number, and corresponds to the *Lighting (degrees)* slider in the dialog.

```
(let* ((width (* radius 3.75))
```

Here, `let*` is a way to set a list of variables in Script-Fu. Variables are quantities that you need to calculate so that you can use them later in the script. This line is setting a variable called *width* (which will be the width of the new image) to 3.75 times the radius chosen in the dialog.

Count the parentheses (counting parentheses is an important part of Script-Fu programming, as you've no doubt noticed by now). By the end of this line, there are two levels still open.

The inner level—the one right after the word `let*`—marks the list of variables. There are going to be quite a few variables defined here, not just *width*, so this parenthesis level is not closed after the *width* line.

The outer level, the one before the word `let*`, defines the *scope* of the variables, or where you can use them. The variables you define inside a `let*` will be available until you close that parenthesis.

```
(img (car (gimp-image-new width height RGB)))
(drawable (car (gimp-layer-new img width height RGB-IMAGE
                             "Sphere Layer" 100 NORMAL-MODE)))
```

In Lisp, you'll see the word “car” a lot. It has nothing to do with hot rods: it means “the first item in a list.” So `(car '(1 2 3 4))` gives you 1 (a single quote means “What follows is a list of constants, not a function call”). Way back in the early days of Lisp, someone thought that “Contents of Address Register” was a good way to express this concept, and now we're stuck with it.

Another term you'll see a lot is “cdr” (short for “Contents of Decrement Register” and pronounced “could-er”), which means “Everything except the first item.” So `(cdr '(1 2 3 4))` is `(2 3 4)`. Sometimes these are combined together: *cadr* is short for the car of the cdr, so `(cadr '(1 2 3 4))` is the same as `(car (cdr '(1 2 3 4)))`, which is (try working this out on your own before reading the solution): 2.

Some Script-Fu routines return a list, but what you actually need to work with is the first element of the list. `gimp-image-new` and `gimp-layer-new` are two such routines, and you'll almost always take the first item, using `car`, and throw away the rest. If you get confused about whether a function needs you to take the car, just look at other scripts and see how they use that function.

```
(gimp-context-push)
```

GIMP has a “context” which includes details like the foreground and background colors, the current brush, pattern, and gradient, and some other details such as font, opacity, paint mode, and palette. `gimp-context-push` saves all these values so that you can restore them easily when the script has finished its work. It’s a good idea to do this if your script will change color, brush, or any similar value.

```
(gimp-image-undo-disable img)
```

This temporarily disables GIMP’s Undo stack. As a result, when the script has finished running, a single Undo will undo the whole script, not each separate operation the script did. This is almost always a good idea for scripts and plug-ins.

```
(gimp-image-add-layer img drawable 0)
```

After a layer is created, it has to be added to the image before it can be used. The final `0` is the layer position, meaning add the layer at the bottom of the layer stack. `-1` can be used to add the layer at the top of the layer stack; any other number specifies a particular layer position.

```
(gimp-context-set-foreground sphere-color)
(gimp-context-set-background bg-color)
(gimp-edit-fill drawable BACKGROUND-FILL)
```

You’ve already seen what these do: set foreground and background colors, and then fill with the background color.

If you remove the `gimp-edit-fill` line and replace it with a `gimp-edit-clear` line, as in the last section, you can also remove the `gimp-context-set-background` line, since we’re about to set the background color again. You could also remove the background color selector from the dialog; we’ll get to that in a second.

```
(gimp-context-set-background '(20 20 20))
```

Colors in Script-Fu are represented as a list of the red, green, and blue values. This sets the background color to `(20, 20, 20)`: a very dark gray. Remember, the single quote at the beginning of `'(20 20 20)` means “What follows is data, not a function call.” Without the quote, Script-Fu would look for a function named `20` and try to run it with parameters `20` and `20`.

```
(if (and
    (or (and (>= light 45) (<= light 75))
        (and (<= light 135) (>= light 105))))
    (= shadow TRUE))
```

This is a conditional—an “if” statement. Unfortunately, in Lisp they’re a bit hard to read when you’re getting started.

Think back to the way parenthetical expressions “nest” inside each other. Terms like “and” and “or” come first, followed by the two expressions being combined. `(<= light 75)` tests whether the light angle is less than or equal to 75, while `(= shadow TRUE)` tests whether shadow (the true/false parameter attached to the checkbox in the dialog) is true.

So this long *if* clause ends up saying: “If the light angle is greater than 45 degrees AND less than 75, OR if it’s between 135 and 105 degrees, AND if we’re supposed to draw a shadow, then . . .” and the next 11 lines are all part of that shadow-drawing clause.

```
(let ((shadow-w (* (* radius 2.5) (cos (+ *pi* radians))))
```

You may remember `let*` earlier in the script, which lets you set temporary variables that you'll use later. This is setting width, height, `x`, and `y` for the shadow.

The only difference between `let*` and `let` is that with `let*`, you can set variables in order (variables set later in the list can refer to earlier variables), and you can't in `let`. So in this case, if `shadow-h` was set to `(shadow-w * 2)`, `let` would give an error, while `let*` would work.

If that confuses you, just use `let*` everywhere instead of `let`.

```
(if (< shadow-w 0)
```

If the shadow width is less than zero,

```
(prog1 (set! shadow-x (+ cx shadow-w))
      (set! shadow-w (- shadow-w))))
```

A Lisp “if” statement ordinarily takes a single clause after it, plus a second one, which is presumed to be an “else”: `(if (condition) (do-something) (do-something-else))` will *do-something* if *condition* is true, otherwise it will *do-something-else*.

`prog1` lets you insert several expressions in an “if” statement (or other statements like it) where normally only one would be allowed. `(prog1 (do_thing_1) (do_thing_2) (do_thing_3))` is treated as all one statement as far as “if” is concerned.

`set!` sets a variable; it's a little like `let*`, except that in `let` and `let*`, the value will only apply inside the `let` or `let*`. With `set!`, you can change a value and have it persist even outside the clause (such as the *if* here) that contains it.

So if the shadow width is negative, set `shadow-x` to be the width of the window minus the shadow, and set `shadow-w` to its absolute value. This will make the shadow go in the opposite direction.

```
(gimp-ellipse-select img shadow-x shadow-y shadow-w shadow-h
  CHANNEL-OP-REPLACE TRUE TRUE 7.5)
(gimp-edit-bucket-fill drawable BG-BUCKET-FILL
  MULTIPLY-MODE 100 0 FALSE 0 0))
```

This actually draws the shadow, by making an ellipse selection of the right size, and then filling it in Multiply mode. `CHANNEL_OP_REPLACE` means “Replace the contents of the current selection,” rather than adding, subtracting, or intersecting with it—equivalent to the Mode tabs in the tool options for Ellipse Select.

```
(gimp-ellipse-select img (- cx radius) (- cy radius)
  (* 2 radius) (* 2 radius)
  CHANNEL-OP-REPLACE TRUE FALSE 0)
```

Select the circle that will become the actual sphere.

We're back outside the shadow-drawing clause now, so this will be run every time, whether or not there's a shadow.

```
(gimp-edit-blend drawable FG-BG-RGB-MODE NORMAL-MODE
  GRADIENT-RADIAL 100 offset REPEAT-NONE FALSE
  FALSE 0 0 TRUE
  light-x light-y light-end-x light-end-y)
```

This is the tricky line that actually makes the shading on the sphere. In fact, it's a simpler method than the one you used in Chapter 9. Rather than filling the whole circle with a color, shading it, and adding a highlight, you can do the whole operation using a single gradient. Most of the rest of the items are just settings in the Gradient tool options. So pull up a GIMP Toolbox, select the Gradient tool, and compare the tool options with the parameters in this call:

- `gimp-edit-blend` is the command to draw a gradient, just like the Gradient tool would do. (You may remember from Chapter 4 that the Gradient tool is also called the Blend tool.)
- `drawable` is the layer where the drawing will happen.
- `FG-BG-RGB-MODE` is the “blend mode” (not to be confused with drawing mode); this is usually the right setting.
- `NORMAL-MODE` is the drawing mode, corresponding to the *Mode* dropdown: `NORMAL` draws on top of anything that's already there.
- `GRADIENT-RADIAL` is the *Shape* option.
- `100` is the *Opacity*.
- `offset` corresponds to the *Offset* slider.
- `REPEAT-NONE` is the *Repeat* setting.
- `FALSE, FALSE, 0, 0, and TRUE` set *reverse*, *supersample*, *max\_depth*, *threshold*, and *dither* to reasonable values.

The remaining four options are the coordinates to use for the blend: basically, they correspond to where you would drag if you were making the gradient interactively.

If you want to try this yourself, make a circular selection and set the foreground and background colors—remember the background color was set earlier to (20, 20, 20)? Then set all the Gradient tool options to mimic what the script is doing, and drag from somewhere near the upper-right of the circle, where the highlight should be, down and left to the edge of the circle. The Gradient tool will do the rest, and should give you a nicely shaded sphere.

```
(gimp-selection-none img)
(gimp-image-undo-enable img)
```

To clean up, remove the circular selection, and re-enable Undo in the image. An Undo after running the script will undo everything back to the `gimp-image-undo-disable` call—in other words, it will undo the whole script.

```
(gimp-display-new img)
```

Display the image. If you don't include this in your script, the new image will never appear.

```
(gimp-context-pop)))
```

More cleanup: restore the context (particularly the foreground and background colors) from before the script was run. Notice the two extra close-parentheses at the end: these close the first `let*` and the definition of the function `script-fu-sphere`.

### Registering the Script

We're done with the routine that draws the script! All that's left is to register it so that it shows up in the menus, and to specify what options the dialog should have.

```
(script-fu-register "script-fu-sphere"
```

`script-fu-register` takes several parameters. You have to get them in the right order. If you write a new script, it may be easiest to copy an existing one and just change all the names. The first parameter (in quotes above) is the name of the function that will do the work for the script.

```
  _" _Sphere..."
```

This is the name that will be shown to the user: it will appear in menus and in the dialog. The underscores have the functions discussed earlier: internationalization and menu mnemonics.

```
  "Simple sphere with a drop shadow"
```

This is a description of what the routine does. This doesn't need an underscore, because it's not normally shown to the user (except in the Plug-In Browser window) and won't normally be translated.

```
  "Spencer Kimball"
  "Spencer Kimball"
  "1996"
```

This is the script author and copyright holder (these two are almost always the same), and the date.

```
  ""
```

This is where the list of valid image types would go for a script that runs from an image window, such as "\*" or "RGB\*". Since the Sphere script makes a new image rather than modifying an existing one, this is blank.

```
  SF-ADJUSTMENT _"Radius (pixels)"      '(100 5 500 1 10 0 1)
  SF-ADJUSTMENT _"Lighting (degrees)"   '(45 0 360 1 10 0 0)
  SF-TOGGLE     _"Shadow"               TRUE
  SF-COLOR      _"Background color"     '(255 255 255)
  SF-COLOR      _"Sphere color"        '(255 0 0)
```

After the image type comes the list of options to be shown in the dialog. These must correspond to the parameters the function `script-fu-sphere` takes. You can use as many or as few as your script needs. See the next section for a list of all the possible user interface options and how to use them.

In your `my-sphere` script, you removed the call that used the background color. It would be a good idea to remove the "Background color" line from this registration routine too, so that the user doesn't get confused wondering why it's asking for a background color it never uses. If you do that, be sure to remove the `bg-color` parameter from the definition of `script-fu-sphere` at the beginning of the script, or GIMP will give you an error when you try to run the script.

```
(script-fu-menu-register "script-fu-sphere"
  "<Toolbox>/Xtns/Misc")
```

The Sphere script uses these lines to register its position in GIMP's menus. Many scripts do this as part of the normal `script-fu-register` call, and don't use `script-fu-menu-register` at all. If Sphere had done that, its register call would have started like this:

```
(script-fu-register "script-fu-sphere"
  "<Toolbox>/Xtns/Misc/Sphere..."
```



However, the GIMP developers recommend that you use the separate `script-fu-menu-register` for new scripts: it offers more flexibility, such as letting a script appear in two different menu locations.

## User Interface Options for Scripts

Scripts—in any language supported by the GIMP—can choose from a variety of interface widgets to let the user set values. In Script-Fu, these are names in all capital letters preceded by “SF-” (for “Script Fu”): for instance, `SF-TOGGLE`. All GIMP plug-in languages have the same set of options, but some use different prefixes instead of “SF-”.

For each element shown in the dialog, the first parameter is a label to be shown to the user: something like “Radius (pixels)” or “Background color.”

### TOGGLE

This can be a true or false value. It takes one parameter, either `TRUE` or `FALSE`.

### VALUE, STRING

This is a single line string value, such as “Hello, world”. You can use either `VALUE` or `STRING` interchangeably.

### TEXT

This is a multi-line text value.

### RADIO

This is a collection of values that are all tied together, like the station buttons on a car radio. Only one can be chosen at a time: when you press a button, whatever was previously selected becomes deselected.

### ADJUSTMENT

This is a way of specifying a number. It can be a minimalist *spin box*, a text field with up and down arrows, like the `radius` parameter in the Sphere script. Or it can be a full-blown slider like the `lighting` parameter in the Sphere script.

The parameters are as follows:

- *value*: The initial value for the number
- *lower*: The smallest value allowed
- *upper*: The largest value allowed
- *step\_inc*: How much the number will change for each click of the arrow buttons, or when using the arrow keys
- *page\_inc*: How much the number will change when clicking in the slider’s “trough” or using the Page Up/Page Down keys

- *digits*: The number of places after the decimal point. Use 0 if you want an integer value.
- *type*: Either SF-SLIDER (0 will also set it to a slider) or SF-SPINNER (1 will work for a spinner).

## COLOR

This is a color. The GIMP will show a button with the default color you specify, and clicking on that button will pop up a color picker.

## FILENAME

This is the name of a file (usually an absolute pathname on your system). It shows a button that lets users pop up a file chooser (to pick another file).

## DIRNAME

This is the name of a directory or folder on your system.

## OPTION

This is a way to show a “combobox” where the user can select from several values. The parameter is a list of string options: for instance, ("Horizontal", "Vertical").

This will be returned to the script as an integer, the number of the value chosen, starting at 0. So in the example given, if the user selects “Horizontal,” the script would see 0; if the user selects “Vertical,” 1.

## ENUM

This is a way to let the user choose between some predefined values within GIMP. *Enum* stands for “enumerated,” because the possible values are specified by number in a fixed list. Parameters are as follows:

- Enum name (e.g., “InterpolationType”)
- Initial value (e.g., “linear”)

## BRUSH

This allows the user to choose a brush. Its parameters are as follows:

- Brush name (e.g., “Circle (03)”)
- Opacity
- Spacing
- Paint Mode, e.g., Normal

## PATTERN

This is the name of a pattern as it appears in the Patterns dialog, e.g., “Pine.”

## GRADIENT

This is the name of a gradient as it appears in the Gradients dialog, e.g., “FG to BG (RGB).”

## PALETTE

This is the name of a palette for indexed colors, e.g., “Web.”

## FONT

This allows the user to choose a font. It takes one parameter, the initial font name, such as “Helvetica.”

## IMAGE, DRAWABLE, LAYER

These are used to pass in the image and drawable used by the plug-in, if any. They take two parameters: usually you’ll pass in a name, and 0 (zero), as the `erase-rows.scm` script did. The system will take care of replacing that 0 with the currently active layer.

Usually, you’ll see `Image` and `Drawable` as the first two parameters in any script that works on an existing image, like `erase-rows`.

## A Python Script

Conveniently, there’s a `Sphere` script in Python, too, and it works exactly the same way as `Script-Fu`. It’s called `sphere.py`, and you’ll find it in your GIMP system plug-ins folder if you have GIMP Python installed. If not, you can get it from the online “CVS” source repository (see Chapter 12 for more details on the CVS viewer): <http://cvs.gnome.org/viewcvs/gimp/plug-ins/pygimp/plug-ins/sphere.py>. Comparing the two scripts is a great way to compare the two scripting languages. You’ll probably find Python quite a bit easier to read than `Script-Fu`.

Pull up `sphere.py` in your favorite text editor. The structure should look very familiar once you’ve been through `sphere.scm`. There’s a routine that actually creates the sphere, and then a routine that registers the script and inserts it into the menus. Let’s walk through `sphere.py` to see the differences:

```
#!/usr/bin/env python
```

This line tells the operating system that the file is a Python script. This special character combination `#!` is called the *shebang*, because `#` is often pronounced “hash” and “bang” by programmers. You’ll see a similar line at the beginning of the Perl example later in this chapter.

```
# Gimp-Python - allows the writing of Gimp plugins in Python.
# Copyright (C) 1997 James Henstridge <james@daa.com.au>
```

In Python, the comment character is a hash (`#`), not a semicolon (`;`). (Yes, that means the “shebang” line is technically a comment.)

```
import math
from gimpfu import *
```

The Python `Sphere` script imports Python’s `math` library because it does some basic trigonometry and uses the value of `pi`. If you need any other specialized Python functions, you can import them here.

All Python GIMP scripts must include the second import line. That's what gives you access to all the GIMP functions. The special syntax of the import statement lets you use GIMP functions without inserting the sequence `gimpfu.` in front of each call.

```
def python_sphere(radius, light, shadow, bg_colour, sphere_colour):
```

`def` is how Python defines routines. Otherwise this is much like Script-Fu.

```
    if radius < 1:
        radius = 1
```

Comparisons and variable assignments are much easier to read in Python than in Script-Fu.

In Python, indentation is important: it shows the structure of the program. A block of code inside an `if` clause must be indented more than the `if` itself; indentation is the only way to indicate that structure in Python, where other languages might use `begin` and `end`, braces (`{and}`), or parentheses. If you copy Python code, be sure to include all the spaces at the beginning of each line so you don't accidentally change indentation levels, and don't use tabs; they can confuse you (and Python) as to the exact indentation level. You're best off using a four-space indent for GIMP code, since that's what the standard GIMP plug-ins use.

```
        width = int(radius * 3.75)
        height = int(radius * 2.5)
```

Calculate the size of the canvas. `int(value)` converts a floating-point number to an integer.

```
        img = gimp.Image(width, height, RGB)
```

Python is "object-oriented." There's an object called `gimp` that contains all the important GIMP functions you'll be using. `gimp.Image` says to create an `Image` object from the object `gimp`. You can use the `Image` object (called `img`) later in your script.

```
        drawable = gimp.Layer(img, "Sphere Layer", width, height,
                               RGB_IMAGE, 100, NORMAL_MODE)
```

A `Layer` is also an object: this creates a new layer.

```
        old_fg = gimp.get_foreground()
        old_bg = gimp.get_background()
```

In addition to making new objects, the `gimp` object has routines like `get_foreground()` to fetch the current foreground color. This script saves the old colors here, so it can restore them at the end. This is the equivalent of the `gimp-context-push` call you saw in Script-Fu. Python can also push and pop the whole context; whether to do that, or simply save a couple of colors, is a matter of preference.

```
        img.disable_undo()
        img.add_layer(drawable, 0)

        gimp.set_foreground(sphere_colour)

        gimp.set_background(bg_colour)
```

Again, this has all the same operations the Script-Fu version did.

```
        pdb.gimp_edit_fill(drawable, BACKGROUND_FILL)
```

Some functions, rather than being in the `gimp` object, are in another object called `pdb`. I'll talk about the *Procedural DataBase* in the next section. Functions in the `pdb` object are generally named `gimp_functionname`.

Notice that Python uses underscores where Script-Fu used dashes. If Python sees a dash, it assumes you want to subtract something from something.

```
gimp.set_background(20, 20, 20)

if (light >= 45 and light <= 75 or light <= 135 and
    light >= 105) and shadow:
```

Python conditionals are much easier to read than the Lisp version. When there are this many options, though, you still have to stop and think about it.

The rest of the script is similar, so I'll skip ahead to the registration routine:

```
register(
    "python_fu_sphere",
```

Notice that unlike Script-Fu, this isn't the name of the routine that actually does the work. That routine is `python_sphere`. This is a name that will be made available externally in case another script wants to use the Python Sphere script.

```
    "Simple spheres with drop shadows",
    "Simple spheres with drop shadows (based on script-fu version)",
```

The description of the routine is repeated twice. The first is the short description; the second is a somewhat longer version to show up in online help.

```
    [
        (PF_INT, "radius", "Radius for sphere", 100),
        (PF_SLIDER, "light", "Light angle", 45, (0,360,1)),
        (PF_TOGGLE, "shadow", "Shadow?", 1),
        (PF_COLOR, "bg_colour", "Background", (255,255,255)),
        (PF_COLOR, "sphere_colour", "Sphere", (255,0,0))
    ],
```

Notice that Python uses square brackets (“[]”) to delimit the list of widgets to be shown in the dialog. Also, it precedes the names of the widgets with “PF\_” instead of “SF\_”—that stands for “Python Fu”—and note again how Python uses underscores instead of dashes.

```
    []),
```

If the `python_sphere` routine returned useful values that could be used in calculations by other scripts, an indication of that would go here. Most interactive scripts just use empty brackets like this.

```
    python_sphere)
```

Finally, at the very end of the `register` function comes the name of the routine that will do the work.

```
main()
```

You need to include this line at the end of any Python GIMP script: this is how GIMP actually runs the script.

## Finding Out How to Do Things: The GIMP PDB

How do you find out the names of all those GIMP routines you can call? How did the author of the Python script know to call `pdb.gimp_edit_fill` in order to run the Bucket Fill tool?

First, you can look in other scripts. Often, that's the best way. If you know a script does something similar to your goal, simply read that script and find out how it works. If you're not sure, try searching for keywords such as “fill” in all the scripts GIMP has installed.

If a quick search through other scripts doesn't turn up what you need, you need the Procedure Browser. *Xtns* ► *Procedure Browser* brings up the window shown in Figure 11-9.

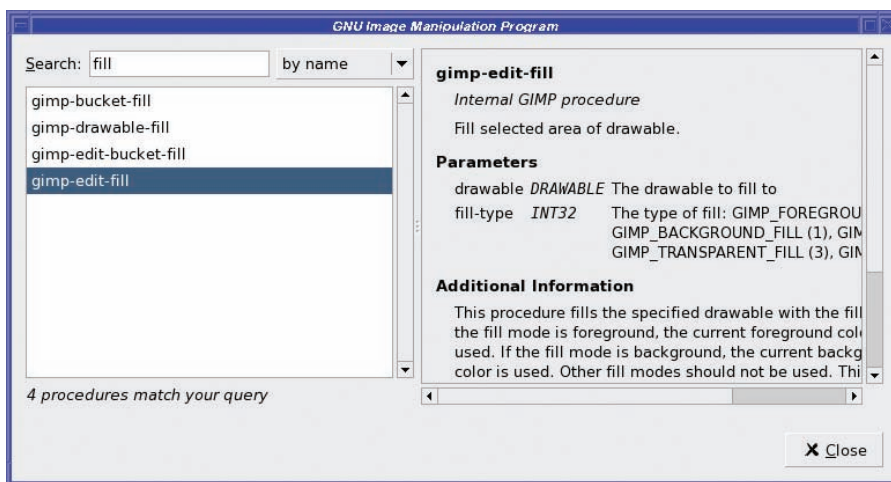


Figure 11-9. The GIMP Procedure (PDB) Browser

It works a lot like the GIMP Plug-In Browser. You can type a search term, such as “fill,” to search for only procedures that contain that term.

In this case, there turn out to be four different fill procedures. Click on each of them in turn and read the help in the right half of the window to find out what they do.

- *gimp-bucket-fill* says “This procedure is deprecated!” So you don't want to use that.
- *gimp-drawable-fill* fills a drawable (such as a layer) with the current foreground or background, white, transparent, or the current pattern.
- *gimp-edit-bucket-fill* works more like the Bucket Fill tool: it includes parameters for all the options that the tool offers, such as paint-mode, opacity, threshold, sample-merged, and so on.
- *gimp-edit-fill* fills a drawable using the same options as *gimp-drawable-fill*, but it also clarifies the difference: *gimp-edit-fill* will fill a selection if there is one, while *gimp-drawable-fill* will always fill the whole layer even if there's a selection present.

The Procedure Browser is an excellent way of finding details on all the useful functions GIMP makes available to scripts and plug-ins. Don't forget it's there.



## A Perl Script

As mentioned earlier, Perl is no longer a supported part of GIMP. But some people still use it, and there are several useful Perl scripts available on the web. If you're a Perl programmer, you will probably want to install `gimp-perl` and try your hand at writing scripts.

GIMP doesn't come with a Perl version of the Sphere script. So instead, take a look at one of the simplest scripts it does have: *Mirrorsplit*, which splits an image in half and replaces one of the halves with a mirror image of the remaining half. If you have `gimp-perl` installed, you'll find this in your system plug-ins directory. Otherwise, you can get the source from GIMP's CVS: <http://cvs.gnome.org/viewcvs/gimp-perl/examples/mirrorsplit>.

Load it in your text editor. Let's take *Mirrorsplit* bit by bit:

```
#!/usr/bin/perl
```

This “shebang” line tells the operating system that this is a Perl script. Perl's comment character is a hash (#) just like Python's.

```
eval 'exec /usr/bin/perl -S $0 ${1+"$@"}'
    if 0; # not running under some shell
```

Perl magic. Just copy it.

```
use Gimp qw(:auto __ N_);
use Gimp::Fu;
use Gimp::Util;
```

These lines tell Perl that it will be using GIMP routines. The odd first line has to do with reading text.

---

**Tip** The command `perldoc Gimp` will give more information on what that first *use* line is doing, as well as other detailed information on GIMP Perl. The capital “G” is important. `perldoc Gimp::Fu` gives a good general introduction to writing GIMP scripts in Perl.

---

```
use strict;
```

This makes Perl stricter about warning you on possible syntax errors. It's not necessary, but it's a good idea.

```
register      "mirror_split",
             "Splits and mirrors half of the image, according to settings.",
             "Just tick appropriate radio button.",
             "Claes G Lindblad <claesg@algonet.se>",
             "Claes G Lindblad <claesg@algonet.se>",
             "990530",
             N_ "<Image>/Filters/Distorts/MirrorSplit...",
```

Like Python, Perl's `register` function lets you specify a string describing the function, and another one for help, followed by date, author, copyright owner, date, and location in the menus.

`N_` in Perl performs much the same function that an underscore alone performed in Script-Fu and Python: it marks strings for translation into other languages.

```
"*",
```

Unlike Sphere, the `Mirrorsplit` script modifies an existing image. So, instead of leaving the image type blank, the script specifies `"*"`, a wildcard that means it can work on any type of image.

```
[
  [PF_RADIO, "mirror", "Which half to mirror?", 0,
    [Upper => 0, Lower => 1, Left => 2, Right => 3]
  ]
],
```

The list of interface widgets to use begins with a single open bracket (`[]`). Then each item—there's only one in this case—is delimited by its own set of brackets, and finally there's a close bracket. Notice that Perl uses a `PF_` prefix, just like Python does.

`"mirror"` is a radio box, a collection of values from which the user can select only one. `Upper`, `Lower`, `Left`, and `Right` are the possible values: the arrows (`=>`) tie each of these values to a number. GIMP will use the number when it runs the script.

```
sub {
```

Instead of defining a named function like Script-Fu and Python do, Perl puts the important code in a subroutine inside `register`. If your Perl script is complicated, though, it's probably best to define another routine, and just call it from here.

```
my ($img, $layer, $mirror) = @_;
my $w = $layer->width();
```

`my` is the way Perl sets local variables, and Perl always puts a dollar sign (`$`) in front of variable names.

The rest of the script is straightforward. If you like Perl, and you've read through the Script-Fu and Python portions of this chapter, you'll probably have no problem following the Perl script. If you're not a Perl programmer already, it's probably safer to stick to a language that comes standard with GIMP.

## Writing a C Plug-In

For most tasks, writing a script is easier than writing a C plug-in, and just as good.

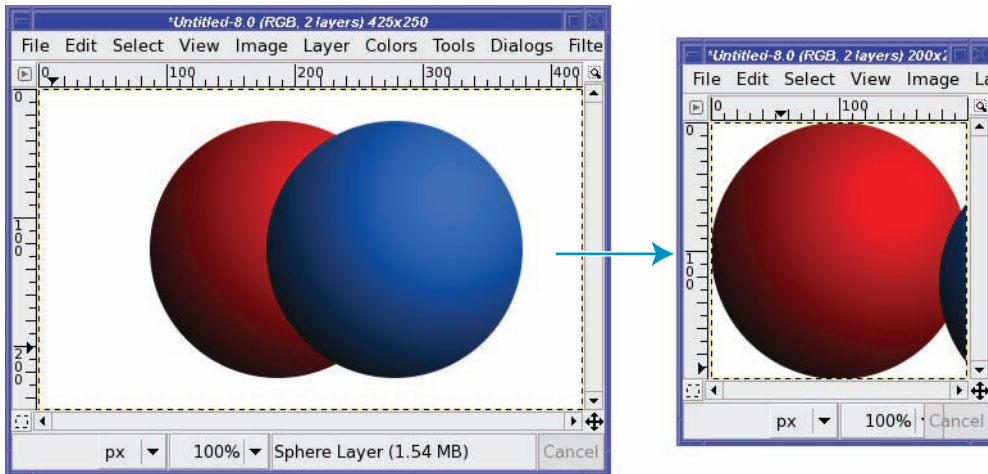
But when you really need performance—especially when you need to loop over each pixel in the image—C is the best bet.

Writing a C plug-in is quite a bit more involved than writing a script, but the basics are the same. Programmers who know a little C—or a related language such as C++ or Java—should be able to create GIMP C plug-ins fairly easily.

To write a C plug-in, you need a development environment installed, including a compiler and the `gimptool` program, as discussed earlier in this chapter.

As with script-writing, it's usually best to learn by modifying an existing plug-in. As an example, I'll start with the *Zealous Crop* plug-in.

Zealous Crop examines the active layer and tries to crop away anything that seems to be background: anything that's transparent or a solid color. That's sometimes a useful function, but if you have multiple layers it probably won't crop the image to the right size. It also does strange things with transparency (Figure 11-10).



**Figure 11-10.** *Zealous Crop doesn't handle multiple layers very well.*

I often want something like Zealous Crop or Autocrop that examines all the layers, not just the active one. Then it would crop the whole image to a reasonable size, including the objects in all the layers. Why not make it happen?

---

**Note** Why start with Zealous Crop rather than Autocrop? Two reasons. First, Autocrop is longer than Zealous Crop, yet the code isn't any easier to read. Second, Zealous Crop illustrates some important principles of C plug-ins which Autocrop does not, so it's a better example for learning plug-in techniques. You could start from either one and come up with a perfectly good Autocrop All Layers plug-in.

---

Start by finding the source of Zealous Crop. If you've already downloaded the GIMP source, you can find it there, in `plug-ins/common/zealouscrop.c`. But if you don't have the GIMP source handy, you can always get it at [cvs.gnome.org](http://cvs.gnome.org), where the GIMP source lives: <http://cvs.gnome.org/viewcvs/gimp/plug-ins/common/zealouscrop.c>.

## A Tour Through Zealous Crop

```
/*
 * ZealousCrop plug-in version 1.00
 * by Adam D. Moss <adam@foxbox.org>
 [ ... ]
 */
```

In C, a comment begins with `/*` and can extend across multiple lines.

The sequence `*/` ends the comment. The asterisks at the beginning of each line in between are simply for readability; C doesn't require them.

C plug-ins typically begin with a comment crediting their author and listing copyright and licensing information. Since the GIMP is free software, licensed under the GNU General Public License, most people who write plug-ins use that license. That way, their plug-ins can be freely shared with other GIMP users (and perhaps, one day, included with GIMP itself, as this one is).

```
#include "config.h"
#include <libgimp/gimp.h>
#include "libgimp/stdplugins-intl.h"
```

`#include` is how C inserts interfaces for other libraries. The two that use quotation marks around the filename are specific to plug-ins built inside the GIMP, and should be removed for external plug-ins. The one for `gimp.h`, with angle brackets, is needed for every C plug-in.

```
#define PLUG_IN_PROC "plug-in-zealouscrop"
```

The name of the plug-in is `plug-in-zealouscrop`, and this name will be used in several places. Rather than duplicate it each time, this lets you use a slightly shorter symbol to refer to it. If you ever change the name of the plug-in (as we're about to do), you only need to change it in one place.

```
GimpPlugInInfo PLUG_IN_INFO =
{
    NULL, /* init_proc */
    NULL, /* quit_proc */
    query, /* query_proc */
    run, /* run_proc */
};
```

The structure of a plug-in includes functions to initialize, quit, and query for information as well as run the plug-in. The scripting languages handle this for you, but in C you must include these lines.

```
static gint bytes;
```

`bytes` is a variable that contains an integer value. `gint` is the way the libraries used underneath GIMP refer to an integer. (The “g” ultimately derives from the GNU project, just like the G in GIMP.)

```
MAIN ()
```

This is a way to let GIMP call the plug-in, much like the similar call at the end of a Python script. Include this unchanged in each plug-in you write. Notice it's capitalized: it's a macro, not the usual C function called `main()`, so that each platform can use the plug-in code in whichever way works best.

```
static inline gint
colours_equal (const guchar *col1,
               const guchar *col2,
               gint          bytes)
```

Now we're getting into the actual programming. `colours_equal` is a C function that returns an integer (`gint`) and takes three parameters: `col1`, `col2`, and `bytes`. `col1` and `col2` are colors, and are specified as an array of `guchar`, unsigned characters (actually a pointer to a character, for those familiar with the C language). `bytes` specifies how many bytes each color includes; this might be 1 for grayscale images, 3 for color images, or 4 for color images with transparency.

```
for (b = 0; b < bytes; b++)
{
    if (col1[b] != col2[b])
        return FALSE;
}
```

Loop over each byte in the two colors and compare them. If they're different at any point, return `FALSE`—the colors are not equal.

```
static GimpParamDef args[] =
{
    { GIMP_PDB_INT32,    "run-mode", "Interactive, non-interactive" },
    { GIMP_PDB_IMAGE,   "image",    "Input image" },
    { GIMP_PDB_DRAWABLE, "drawable", "Input drawable" }
};
```

This describes the parameters the plug-in takes—`args` is short for “arguments,” another word for “parameters.”

For “run-mode” you'll almost always want to use “Interactive” (meaning it may pop up a dialog and take input from the user, and it shows a result). It's possible, though, to write non-interactive plug-ins that are primarily for use by other plug-ins, or a plug-in which needs no parameters passed to it.

This plug-in operates on an input image and drawable, but doesn't take any other parameters.

---

**Note** Why all the emphasis on “drawables”? Why not just say “layer”? Because you can draw on things that are not layers, such as layer masks, channels, or even the QuickMask or a selection.

---

```
gimp_install_procedure (PLUG_IN_PROC,
                       "Automagically crops unused space from "
                       "the edges and middle of a picture.",
```

`gimp_install_procedure` and `gimp_plugin_menu_register` are similar to the `register` routine you've seen in the other languages, and take the same arguments.

## The run Procedure

```
static void
run (const gchar      *name,
     gint             n_params,
     const GimpParam  *param,
```

```

    gint          *nreturn_vals,
    GimpParam     **return_vals)
{

```

Finally, the run procedure. This is the routine GIMP calls when the user actually runs the plug-in.

run procedures in GIMP C plug-ins always look longer and more complicated than you expect, especially if you're used to the more compact form of scripting languages. But a lot of this structure is very standardized, and you can just copy it directly from one plug-in to the next.

```

    GimpPDBStatusType status = GIMP_PDB_SUCCESS;

```

Notice that `status` is initialized to `GIMP_PDB_SUCCESS` here. Later on, that value will be returned, unless an error condition is encountered and `status` is changed to be some other value.

```

    INIT_I18N();

```

This is a GIMP internal function related to internationalization, and you will need to remove it in your own plug-ins.

```

    if (run_mode == GIMP_RUN_NONINTERACTIVE)
    {
        if (n_params != 3)

```

Check the number of parameters. The number must match the number of items in `GimpParamDef`, defined earlier.

This early code is pretty much the same for every plug-in, but after this point we go back to code that might need to be changed.

```

        /* Get the specified drawable */
        drawable = gimp_drawable_get(param[2].data.d_drawable);
        image_id = param[1].data.d_image;

```

C plug-ins need to access their parameters in the proper order before they can reference them by name. In this case, the image is the first parameter of the plug-in, while the drawable is the second parameter. (For some reason, the plug-in author has written the line to get the second parameter `param[2]` before the first parameter `param[1]`. What's important is the parameter numbers, `[2]` and `[1]`, not the order in which they're accessed in the C code.)

These two lines fetch the data part (`data`) of these two parameters, convert them to the right type (`d_image` or `d_drawable`), and then assign them to variables `drawable` and `image_id` so they can be used more easily.

```

        /* Make sure that the drawable is gray or RGB or indexed */
        if (gimp_drawable_is_rgb (drawable->drawable_id) ||
            gimp_drawable_is_gray (drawable->drawable_id) ||
            gimp_drawable_is_indexed (drawable->drawable_id))

```

Checking the image type is a good idea, especially if your plug-in can only work on certain types. In this case, though, it might be a bit excessive: what could the image be if not RGB, gray, or indexed? (Yes, these tests do work even if the drawable has an alpha channel.)

```

        gimp_progress_init (_("ZealousCropping(tm)..."));

```



You can display messages in GIMP's status bar during a plug-in, and display a progress bar to show how far along the plug-in is (you'll see that in a moment). Any type of plug-in can update GIMP's progress bar, including scripts, though the scripts shown earlier in this chapter have been fast-running, so they didn't need that.

```
gimp_tile_cache_ntiles (1 +
    2 * (drawable->width > drawable->height ?
        (drawable->width / gimp_tile_width()) :
        (drawable->height / gimp_tile_height())));
```

C (and also Python) plug-ins have a speed secret compared to Script-Fu: they can use *tiling* to access pixels in the image. Tiles are a way of grabbing a batch of pixels into memory all at once. Plug-ins can then loop through the smaller batch instead of requesting the pixels at specific X and Y coordinates one after another. I'll explain more about tiling in a bit.

```
do_zcrop(drawable, image_id);
```

Call a function (defined below) which does the crop.

```
if (run_mode != GIMP_RUN_NONINTERACTIVE)
    gimp_displays_flush ();
```

The crop is finished (do\_zcrop did all the rest of the work). This updates GIMP's display if it's been run in interactive mode.

```
gimp_drawable_detach (drawable);
```

When a C plug-in is finished working with a drawable, it should “detach” from it, which ensures that all tiles have been saved properly and all extra memory freed.

## The Routine That Does the Work

```
static void
do_zcrop (GimpDrawable *drawable,
          gint32      image_id)
```

Now the do\_zcrop function, which does most of the work, begins. It takes two parameters: a drawable and an image. The first few lines of the function merely declare local variables.

```
width = drawable->width;
height = drawable->height;
```

Get the width and height of the drawable.

If you haven't forgotten the goal of modifying the plug-in to crop images with multiple layers, this may grab your attention. Getting the size of just the current layer won't be enough. Perhaps you'll have to modify this to be the size of the whole image, not the drawable. Or maybe this function should be called during a loop over all layers in the image. I'll examine those options shortly.

```
bytes = drawable->bpp;
```

drawable->bpp is the “bits per pixel” of the drawable: probably 24 for an RGB image, 32 if it has an alpha channel.

```
killrows = g_new (gint8, height);
killcols = g_new (gint8, width);
buffer = g_malloc ((width > height ? width : height) * bytes);
```

`g_new` and `g_malloc` are two ways of allocating memory. `killrows` and `killcols` end up being arrays of integers big enough to hold one color channel (a `gint8` is an 8-bit integer, just big enough for a single-color channel) in one row or one column. `buffer` is large enough to hold all three- (or four-) color channels, and is big enough to hold either one row or one column—it takes the maximum of the two dimensions.

```
/* initialize the pixel regions */
gimp_pixel_rgn_init (&srcPR, drawable, 0, 0, width, height, FALSE, FALSE);
gimp_pixel_rgn_init (&destPR, drawable, 0, 0, width, height, TRUE, TRUE);

livingrows = 0;
for (y = 0; y < height; y++)
{
    gimp_pixel_rgn_get_row (&srcPR, buffer, 0, y, width);
}
```

Loop over tiles (pixel regions) in order to read the individual pixels in the drawable.

## Getting Pixels: Tiles and Pixel Regions

To speed up its operation, GIMP internally accesses drawables one small “tile” at a time. These are usually  $64 \times 64$  pixels, but don’t count on that being the case. Most access to pixels inside the files is done through *pixel regions*. (Python can also access pixel regions.)

### Using Pixel Regions

A pixel region (type `GimpPixelRgn` in C) represents some area of an image that you want to access. It may be a row or column, a rectangle of any size, or a single pixel. Initialize it like this:

```
void
gimp_pixel_rgn_init (GimpPixelRgn *pr,
                    GimpDrawable *drawable,
                    int x, int y, int width, int height,
                    int dirty, int shadow);
```

Here, the parameters are as follows:

- `pr`: This is your local `GimpPixelRgn` variable.
- `drawable`: This is the drawable to be used.
- `x, t, width, height`: This is the area of the drawable you want to access.
- `dirty`: Set this to 1 if you’re planning to write back to the drawable. If you’re only reading it, you can set it to 0.
- `shadow`: If set to 1, GIMP will provide a place to write your new data. Set it to the same value as `dirty`.

Once the pixel region is initialized, you can fill it with pixels from the image by calling functions such as `gimp_pixel_rgn_get_pixel` (to get a single pixel), `gimp_pixel_rgn_get_row`, `gimp_pixel_rgn_get_col`, and `gimp_pixel_rgn_get_rect` (to get an arbitrary rectangle from the image, specified by `x`, `y`, `width`, and `height`). For each of these routines, you need to pass in a pointer to a region of allocated memory large enough to hold the data. In this case, where the plug-in reads one row or column at a time, that is:

```
buffer = g_malloc ((width > height ? width : height) * bytes);
```

The memory buffer's size, in bytes, is the maximum of `width` and `height` (so it can hold either a row or a column) times the number of bytes per pixel (`drawable->bpp`).

## Reading and Writing the Memory

How do you access the pixels after calling `gimp_pixel_rgn_get_row`?

The pixels in the buffer begin at the location you asked for. So in this case, `buffer` will contain a single (horizontal) row. Within that row, each pixel occupies as many bytes as there are color channels. So in the case of an RGBA image (4 bytes per pixel), the first pixel starts at the beginning of the buffer (location 0, since C always begins arrays at 0), the second starts at location 4, the third at 8, and so forth.

That's fairly straightforward for a single row spanning the width of the whole image. But it gets a bit more complicated when you're calling `gimp_pixel_rgn_get_rect` to get a rectangle from somewhere in the middle of the image. In this case, you have to keep track of the location where you got the rectangle, and its size.

You also need to know that the data in pixel regions is organized by row, not by column. Looping through a pixel region, you will see all the pixels for one row, and then immediately afterward you'll see the first pixel from the next row. In other words, X coordinates increase faster than Y coordinates.

For instance, to get the pixel at `x1`, `y1`, with a bytes per pixel of `bpp`, you would read the bytes starting at  $(y1 - y) * width * bpp + (x1 - x) * bpp$ . Got that? The first part,  $(y1 - y) * width * bpp$ , gets you from the top (row `y`) down to row `y1`: each row takes up  $width * bpp$  bytes, and there are  $(y1 - y)$  rows since the beginning of this pixel region. The second part,  $(x1 - x) * bpp$ , is how far over in the row we need to be to get to `x1`. Whew!

## Tiles and the Tile Cache

Using pixel regions, especially large rectangular regions, improves performance quite a lot versus getting a single pixel at a time. But you can improve performance even more by working directly with the tiles GIMP uses. Normally, each tile is passed as a block of shared memory from the application to the plug-in. This can cause performance delays if the plug-in switches tiles often. In that case, you can keep a cache of tiles in the plug-in by specifying the number of tiles to be cached.

You can use this call in combination with pixel regions to make your plug-in faster. Say you're working primarily with rows or columns, as *Zealous Crop* does. Then it helps to cache enough tiles to span an entire row or column. If you'll be writing back to the `drawable`, cache twice as many tiles: you need one set for reading, and another set ("shadow" tiles) for writing.

You may remember from *Zealous Crop*'s run procedure:

```
gimp_tile_cache_ntiles (1 +
    2 * (drawable->width > drawable->height ?
        (drawable->width / gimp_tile_width()) :
        (drawable->height / gimp_tile_height())));
```

This is taking the maximum of the drawable's width and height, dividing by the width or height of GIMP's tiles (usually 64, but you shouldn't depend on it: always use `gimp_tile_width()` and `gimp_tile_height()`), and multiply by 2 to account for the shadow tiles.

Why add 1 to the result? Because integer math (like those divisions) rounds down to the nearest integer, so if you don't add 1, you may end up with a fraction of a tile less than you actually need.

### SPEEDING THINGS UP BY TAKING PIXEL REGIONS DIRECTLY FROM TILES

Another way to use tiles to help make your pixel-region code run faster is to use `gimp_pixel_rgns_process` to take pixel regions directly from GIMP's tiles. Used instead of routines like `gimp_pixel_rgn_get_row` or `gimp_pixel_rgn_get_rect`, this is a routine that automatically gets sequential chunks of data that are aligned to GIMP's own tiles. This avoids a lot of extra copying of pixels, and can be much faster.

If you're writing a plug-in that needs that kind of speed, be sure to check out the plug-in tutorial on [developer.gimp.org](http://developer.gimp.org), which includes examples of pixel-region processing code.

## Searching the Image for Blank Rows and Columns

When we last left our hero (the Zealous Crop plug-in), it had just requested a pixel region row. Now you should have a better idea what that means. (Don't worry if it isn't completely clear. Working with pixel regions is confusing to most people, and it takes time to get used to thinking in those terms.)

```
livingrows = 0;
for (y = 0; y < height; y++)
{
    gimp_pixel_rgn_get_row (&srcPR, buffer, 0, y, width);

    killrows[y] = TRUE;
```

We don't know yet what `livingrows` is, but we're about to find out. This loop will get each horizontal row of the image, starting at the top (`y=0`) and ending at the bottom. `killrows` is an array that has one entry for each row in the image; this row's entry is initialized to `TRUE`.

```
for (x = 0; x < width * bytes; x += bytes)
```

Loop over each pixel in the row. Remember, `bytes` is the number of bytes per pixel, so each pixel spans `bytes` spaces in `buffer`.

```
if (!colours_equal (buffer, &buffer[x], bytes))
{
    livingrows++;
    killrows[y] = FALSE;
    break;
}
```

Aha! Remember the `colours_equal` routine from earlier? If the color of the current pixel is different from the color of the first pixel in the row, then increment `livingrows` and set `killrows[y]` to `FALSE`.

In other words, if this row isn't all the same color, then this row can't be just background, and shouldn't be cropped away.

```
area += width;
if (y % 20)
    gimp_progress_update ((double) area / (double) total_area);
```

And now you see what `area` was for. It's for updating GIMP's progress bar. `gimp_progress_update` takes a floating point number between 0 and 1. Since the plug-in loops over both rows and columns, it needs some measure of how far along it is in its task. `Area` is a good way of measuring that.

What's the `(y % 20)`? That's intended to ensure that GIMP won't waste too much time updating the progress bar instead of doing calculations. It only updates if `y` has a remainder when divided by twenty.

This seems a bit odd: it will update nineteen out of twenty times. You have to wonder whether the intent was to update only if `(y % 20 == 0)`, which would update only every twentieth row. (We'll test that out in a moment.)

The next clause (omitted here) does exactly the same thing, but looping horizontally over each column in the image.

## Writing to Pixel Regions

The next section performs the crop, removing rows and columns marked "killed." It illustrates how to write to an image using pixel regions.

```
destrow = 0;

for (y = 0; y < height; y++)
{
    if (!killrows[y])
    {
        gimp_pixel_rgn_get_row (&srcPR, buffer, 0, y, width);
        gimp_pixel_rgn_set_row (&destPR, buffer, 0, destrow, width);
        destrow++;
    }
}
```

Looping vertically over rows starting from the top, if a row isn't on the kill list, the whole row is copied directly from the source pixel region into the destination. `destrow` is a count of the number of rows that were copied. Notice that this will happily just skip over any blank (killed) rows and continue on—so killed rows can be omitted from the middle of the image.

```
destcol = 0;
gimp_pixel_rgn_init(&srcPR, drawable, 0, 0, width, height,
                   FALSE, TRUE);
```

The source pixel region is initialized again. Why? Because we've just written those rows back to the image, so the image has changed since the last time we initialized.

With the new source region, we can loop over columns and write only the ones that aren't on the kill list, looping over `x` the same way we just looped over `y`.

Then we're done! The rest is just cleanup.

```
g_free (buffer);
g_free (killrows);
g_free (killcols);
```

All the memory that had been allocated needs to be freed. Otherwise your plug-in will have a “memory leak” and will get bigger and bigger. For a GIMP plug-in, this usually isn't critical since a plug-in runs as a separate process, and the process exits as soon as the plug-in is finished running. But free your memory anyway.

```
gimp_progress_update (1.00);

gimp_image_undo_group_start (image_id);
```

Curiously, the plug-in doesn't start Undo until here. That's because the pixel region operations it's been doing don't affect GIMP's Undo system.

```
gimp_drawable_flush (drawable);
gimp_drawable_merge_shadow (drawable->drawable_id, TRUE);
```

These two lines make sure all the data from the destination pixel region are written back to the image.

```
gimp_image_crop (image_id, livingcols, livingrows, 0, 0);
```

Since the image is probably smaller now, this gets rid of the excess.

```
gimp_image_undo_group_end (image_id);
```

End the undo group.

## Making a New Plug-In

Now that you're an expert on how the Zealous Crop plug-in works, it's time to adapt it to the new Autocrop All Layers function.

The first step is to make a copy of the source file (I'll call it `autocropall.c`).

Next, choose a name for your new function, and edit the C file to change the old name to your new name everywhere. Don't forget to change the author, year, and the name of the plug-in inside the `gimp_install_procedure`.

If you're starting with a plug-in from inside the GIMP source, like Zealous Crop, you'll also have to remove the internationalization routines. Those reference routines inside the GIMP that aren't available to outside plug-ins. That means removing `_()`, `N_()`, and the `INIT_I18N` line. You'll also need to remove the internal `#include` lines. If you don't do that, `gimptool` won't be able to compile your plug-in.

Now you're ready to try compiling. `gimptool --build autocropall.c` builds the plug-in but doesn't install it. `gimptool --install autocropall.c` builds the plug-in and then installs it in your GIMP profile folder (assuming it built without errors, of course). Notice that there are two dashes.

The first time you compile and install your new plug-in, you'll have to restart GIMP before it will show up in the menus. Unfortunately, there's no equivalent to *Refresh scripts* to make new C (or Python) plug-ins register themselves.



Fortunately, once the plug-in has registered, you won't need to restart GIMP again. You can rebuild the plug-in and re-run it, and your changes will take effect. You only need to restart GIMP for a plug-in when you've changed the registration routine somehow: if this is the first time you've built this particular plug-in, if you're changing its name, or if you're changing the parameters it takes.

After restarting GIMP, check your menus. The new plug-in name (such as *Image* ► *Autocrop All Layers*) should appear in the menus now. Try running it. You haven't changed anything related to its function, so it should still do a normal Zealous Crop.

Just to make sure you're really seeing the new plug-in, try a tweak. Remember I mentioned that it looked like the status bar was updating too often, and the  $(y \% 20)$  should probably be  $(y \% 20 == 0)$ ?

Try making that change. Everywhere you see  $\% 20$ , add an  $== 0$  after it. Then run `gimptool --install autocropall.c` to reinstall the new plug-in. Success! On my  $500 \times 400$  test image, the time needed for a Zealous Crop shrunk from 6 seconds to about 2. The status bar update is a little jumpy, but it's worth the time savings.

---

**Note** If you find a bug like this, please report it. I reported this problem, and my change will be in GIMP 2.4. See Chapter 12 for information on reporting bugs.

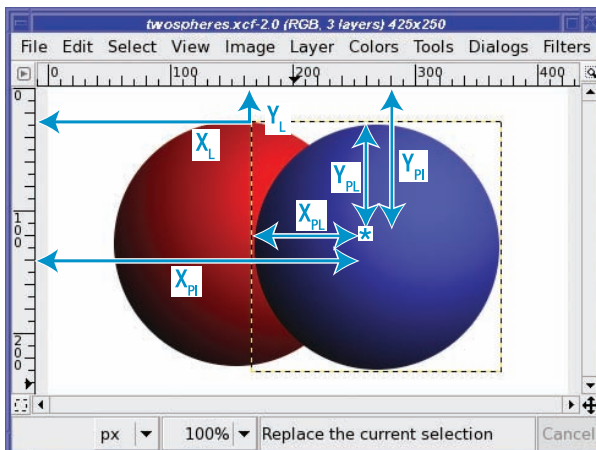
---

## Layer Coordinates vs. Image Coordinates

Now it's time to make the real changes: to make `autocropall.c` look through all the visible layers. The changes will all be in `do_acrop`.

One difference when working with layers, compared to a whole image, is that coordinates may change. Layers aren't always the size of the whole image, and sometimes they're offset compared to the image.

Take a look at the “two spheres” image in Figure 11-11. The layer containing the blue sphere is offset to the right and down compared to the image: the layer's upper-left edge is at image position  $(X_L, Y_L)$ .



**Figure 11-11.** Two spheres, three layers, three coordinate systems

If a point in the blue sphere is at image coordinates  $(X_{PI}, Y_{PI})$ , the point's coordinates within the layer might be much smaller:  $(X_{PL}, Y_{PL})$ . These points are related by simple addition:

$$X_{PI} = X_L + X_{PL}$$

The coordinate of the point in the image is the coordinate of the layer plus the coordinate of the point in the layer, and likewise for the Y coordinates. If you get confused thinking about coordinate systems, always draw a picture—it helps.

Which coordinates did Zealous Crop use? It did everything relative to the drawable—the layer. So it was always using layer coordinates: the  $(X_{PL}, Y_{PL})$  of Figure 11-11.

But wait! Could that be why it gets such strange results for images with several layers, as you saw back in Figure 11-10?

In fact, that's exactly the problem. Try Zealous Crop on your own multi-layered image, and pay attention to where the layer boundaries begin and end. In the final image, the movement of the layers can be explained by where their layer boundaries start.

So for Autocrop All, in addition to looping over all layers, the coordinates of pixels will also need to be handled differently.

## Setting Boundaries

Since Autocrop All doesn't need to crop in the middle of the image, it won't need the arrays `killrows` and `killcols`. All it needs is four integers representing the boundaries where the new crop will take place: call them `xmin`, `xmax`, `ymin`, and `ymax`. `xmax` is initialized to 1 (the left edge of the image) while `xmin` is initialized to `iwidth - 1` (the farthest right pixel), and similarly for `ymax` and `ymin`. So the `killrows` and `killcols` declarations can be replaced with:

```
gint      xmin, xmax, ymin, ymax;
```

`livingrows` and `livingcols` can go away too: we can figure out how many columns will “live” by subtracting `xmax - xmin`, and likewise for rows.

Right after buffer allocation, the plug-in will need to loop over the image layers. Get the number of layers in the image like this:

```
gint      *layers = NULL;
gint      numlayers;
```

```
layers = gimp_image_get_layers (image_id, &numlayers);
```

This will set `layers` to an array of all the layers in the image, and `numlayers` to the number of layers (the number of items in that array).

---

**Tip** How do you find out about calls like `gimp_image_get_layers`? Calls like that aren't in the PDB browser. Happily, *gimp.org* has a very complete programmer's interface manual online, where you can find details of all the calls needed for plug-ins. See the “Additional Resources” section of Chapter 12 for more details.

---

`gimp_image_get_layers` allocates memory, so you should free it when you're finished with it, using `g_free(layers)`.

Instead of `width` and `height` of the drawable, you'll need the width and height of the image (which you can get with `gimp_image_width(image_id)` and `gimp_image_height(image_id)`). Inside the loop over the layers, you'll need the width and height of each drawable. To keep from getting confused about which is which, I've used `iwidth` and `iheight` for the image dimensions, and `dwidth` and `dheight` (declared inside the layers loop) for the drawable dimensions.

`layers` is an array of integers, which represent layer IDs. You can get the drawable for each of these layers with `gimp_drawable_get(layer_id)`. Then you can do all the operations that Zealous Crop did on its drawable, including the pixel-region operations.

## Handling the Edges One at a Time

Once everything is set up, rather than looping over the whole width and height of each layer's drawable, you need only explore inward from each of the four image edges. Keep track of the maximum extent of the crop rectangle so far, in image coordinates, using `xmin`, `xmax`, `ymin`, and `ymax`. Since you aren't looping over the whole drawable, this is very fast; most of the time you'll only need to check a few lines before you find something that can't be cropped. Then you can move on to the next edge.

Since the plug-in will no longer loop over the entire image, there's no point in keeping track of area to update the progress bar. It should be sufficient to update the progress bar once after each layer. Update it to the current layer number divided by the total number of layers. That means you can remove `area` and `total_area` from the program.

Here is how to find `ymax`, the bottom edge from the current layer. The other three edges are similar:

```
start = dheight - 1;
if (layerOffsetY + dheight > iheight)
    start = iheight - layerOffsetY - 1;
```

Find the starting point for the loop. The loop should start at the bottom edge of the image, or the bottom edge of the layer, whichever is higher. (Remember, a layer can run off the edge of the image containing it.)

`dheight` is the height of the drawable (the current layer). Since coordinates start at 0, `dheight - 1` is the coordinate of the bottom-most pixels in the layer.

`layerOffsetY + dheight` represents the image coordinates of the bottom edge of the layer. If that is greater than the image height, then the image's bottom edge lies off the bottom of the image, and `start` is adjusted accordingly, to the last *visible* pixel in the layer.

```
for (y = start; y > 0 && layerOffsetY + y > ymax; y--)
```

Loop backwards (with `y` in layer coordinates) from the bottom visible pixel in the layer until we either reach the top of the image, or until the image coordinates of the current row (`layerOffsetY + y`) crosses the `ymax` we've already found from some other layer.

```
gimp_pixel_rgn_get_row (&srcPR, buffer, 0, y, dwidth);
```

```
for (x = 0; x < dwidth * bytes; x += bytes)
```

Get the row at this position, and loop horizontally over the pixels in it. Notice that this isn't checking to see if the pixels are really inside the image boundaries. That's something that should probably be added.

```
if (!colours_equal (buffer, &buffer[x], bytes))
```

If the color is different from the color of the first pixel in the row,

```
ymax = y + layerOffsetY + 1;
break;
```

convert back to image coordinates and update ymax.

After checking all four edges like this, in each layer in the image, xmin, xmax, ymin and ymax should hold the coordinates to pass to crop. Crop takes its parameters in a slightly unintuitive order: `gimp_image_crop(image_id, new_width, new_height, offset_y, offset_x)`. xmax and ymax are coordinates, not dimensions, so you have to subtract:

```
gimp_image_crop(image_id,
                xmax - xmin, ymax - ymin,
                xmin, ymin);
```

That saves a lot of work compared to writing to pixel regions and updating shadow tiles!

---

**Note** Since you're no longer writing to the image using pixel regions, you can eliminate `destPR` and every line that refers to it. But even better, you can change the `gimp_tile_cache_ntiles` call to remove the factor of two: you only need half the tile cache. That makes the plug-in use less memory.

---

All done! The new plug-in correctly crops multiple layers. Space doesn't permit a full listing here, but you can download the full source for the Autocrop All plug-in, `autocropall.c`, in the Source Code area of the Apress website (<http://www.apress.com>).

## Summary

If you've made it this far, you're well on your way to becoming a GIMP plug-in writer—and if you choose, perhaps a developer for the GIMP itself.

You can develop your plug-in authoring skills by exploring the existing scripts and plug-ins, reading the documentation available on the GIMP website, and writing your own new plug-ins. Whether you're tweaking existing scripts, writing quick scripts for your own use, or developing polished plug-ins intended for distribution, don't be afraid to explore and try new things!

More than any other part of the GIMP, plug-in writing benefits from good documentation. This book isn't intended to be a complete reference manual on GIMP programming; it only scratches the surface. You'll want to become very familiar with the online documentation and existing tutorials.

If you're wondering where to find all these web resources, keep reading. Chapter 12 will discuss the wealth of GIMP information on the web, as well as a few other topics such as GIMP's configuration files, screen shots, and printing.



# Additional Topics

**B**y now, you're familiar with most of the basic techniques of image manipulation available in the GIMP. In this chapter, I'll show you a few details that haven't been covered already, discuss ways of customizing the GIMP, and describe some of the web resources that can help you learn even more. The chapter will cover the following:

- Printing with Gimp-Print/Gutenprint
- Screen shots and scanners: the *Acquire* menu
- GIMP configuration files
- Additional resources

## Printing with Gimp-Print or Gutenprint

Printing in GIMP isn't built-in; it's handled by a plug-in. That means that printing may be different on different platforms.

GIMP for Windows prints using the standard Windows print dialogs. Any adjustment to the output will probably go through the dialogs supplied by your printer manufacturer. These may be different for each printer, so this book won't try to cover them.

On most Linux and Mac systems, the print plug-in for GIMP is included in the normal GIMP installation. It comes from a separate project, though, historically called *Gimp-Print*. (The Gimp-Print project is also the origin of many color printer drivers used by Linux and Mac OS X.) Despite the name, the Gimp-Print project is separate from GIMP and is run by different people. The project was recently renamed Gutenprint, and has a brand-new printing plug-in available. As of this book's press deadline, GIMP still includes the old Gimp-Print plug-in, but it may switch to the new Gutenprint version in the future.

Initiate printing using the menu item *File* ► *Print*. If your image has multiple layers or other GIMP-specific enhancements, GIMP will tell you it needs to “export” the image. Don't worry; this won't change your original image.

Then, you'll see the dialog shown in Figure 12-1.

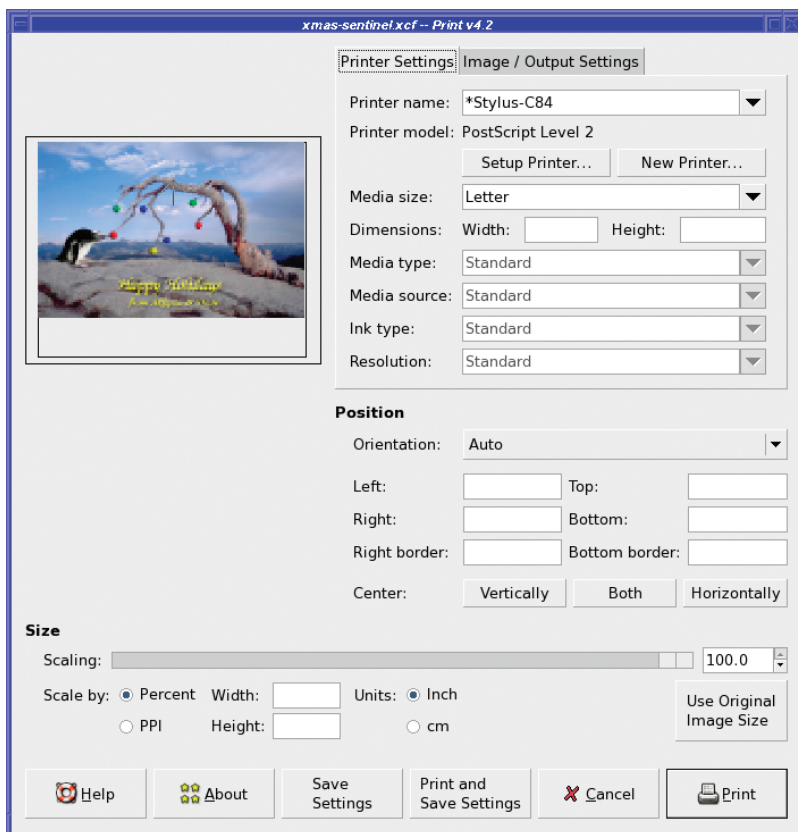


Figure 12-1. The basic Gimp-Print dialog

On the left is a rough preview of your image and how it will look on the page. You can get a general idea of image brightness, but the preview's main use is to let you drag the image to control its size and position on the page.

To the right of the preview is a Settings area with two tabs. Initially, the *Printer Settings* tab is visible. This lets you choose your printer model and control important details like paper size, paper type, and resolution.

## Setting Your Printer Model

The first step is to tell Gimp-Print what type of printer you have by clicking *Setup Printer....* This is **very important**. The single most common cause of printing problems is that Gimp-Print doesn't automatically set this field, even if the printer is already configured as the default system printer.

Notice the printer settings in Figure 12-1. *Printer name* is “\*Stylus-C84.” If I didn't pay close attention, I might be lured into thinking that Gimp-Print knows what type of printer I have. But look at the next line below that: the *Printer model* is “PostScript Level 2.”



Why does it default to PostScript? Because the printing system on Mac OS X, Linux, and most Unix systems is based on Adobe's PostScript language. If GIMP creates PostScript, it can send it to the printing system, which will then translate it to whatever your specific printer needs.

Sounds nice and simple, right? But the generic PostScript model used by the printing system doesn't do high quality printing. It can't control any of the nifty special features most inkjet printers offer, such as high resolution or settings for special photo paper.

For now, try to use a more specific printer model instead of PostScript. Check the *Printer model* and if it's not what you expect to see, click *Setup Printer...* (if there's already a printer displayed) or *New Printer...* (if no printer is displayed yet).

Scroll down to find and select your printer make and model from the list in the dialog. (If you have trouble working in such a small scrolled list, you can resize the dialog larger to make it show more printers at once.) If you don't find your exact printer model, something close *may* work (for instance, C84 works for an Epson C86).

If you don't find anything that's even close, you may have to use the generic *PostScript Level 2* setting. (In the Gutenprint plug-in, this is found under *Printer Make: Adobe*.)

---

**Caution** If you use the PostScript setting, you also need to look at the *Command* field. If it includes “-oraw,” delete that. It's a bug in some versions of Gimp-Print that “-oraw” is included for PostScript printers: it's needed for all other printers, but should not be included for PostScript.

---

If you have a real PostScript printer, the *PPD File* field is a place where you can select a driver for that printer. PPD files apply only to PostScript printers. The field should not appear, or should be grayed out, for everyone else.

Click *OK* and you're back in the Gimp-Print dialog, with the correct *Printer Model*. It's a good idea to click *Save Settings* at the bottom of the dialog immediately after setting your printer, to ensure that GIMP remembers it.

## Set Print Size, Paper Type, and Resolution

Once your printer is set correctly, Gimp-Print's options should reflect the capabilities of your printer. Adjust *Media size* to the size of the paper you're using (Letter, A4, 4 × 6, and so on). For most printers, you shouldn't need to change *Dimensions*, *Media source*, or *Ink type*.

However, *Media type* is important: if you're using special paper, such as Photo Paper or Glossy Film, set it here. You'll find that the paper-type setting makes a big difference in printing photographs: photo quality papers require more ink and can accept higher resolution. Usually, the right paper setting will be obvious, but for some papers, you may have to guess. (Is Premium Matte Photo Paper most like Photo Paper, Matte Paper Heavyweight, Premium Luster Photo Paper, or what? You may have to experiment to get the best results.)

*Resolution* is the final setting, and it should offer a list of the resolutions your printer can use. Higher resolutions will take much longer to print and will use more memory and more ink, but will give you a print with sharper detail and smoother color.

What about that other tab, *Image/Output Settings*? I'll get to that in a minute. But first, take a look at the bottom half of the dialog, starting with that *Position* section on the right below the printer settings area.


## Setting the Image's Position, Orientation, and Size

*Orientation* lets you choose the orientation of the image on the page (Figure 12-2).

Landscape:



Portrait:



**Figure 12-2.** Orientation options

The choices are as follows:

- *Auto*: GIMP will use the orientation that best matches the image's own dimensions.
- *Portrait*: The image will be right-side up when the paper is held like a normal page, taller than it is wide, even if the image is wider than it is high.
- *Landscape*: The image will be right-side up when the paper is held with its longer dimension horizontal.
- *Upside down*: Portrait, but with the image upside down. The bottom of the image will print first.
- *Seascape*: Landscape, but upside down.

Now skip down to the broad *Size* slider. Grab it and slide it to change the size of the image. 100% actually means “the largest size that will fit on the page,” but you can scale the image down as small as you want. The preview will update as you adjust the slider, so you can see how the image compares to the paper size. The *Width* and *Height* of the image as it will be printed is displayed under the slider; if you want the image to come out a specific size, you can edit these fields directly rather than using the slider. But be careful: they don't always update when you move from one to another, so it may work better to use the slider until the *Width* and *Height* read your desired amounts.

The checkboxes under the slider let you change the units, for instance to *PPI* (Pixels per Inch) instead of *Percent*.

As soon as you make any changes to the *Size* slider, you'll see the fields under *Position* update to show the current settings for the *Left*, *Right*, *Top*, and *Bottom* edges of the image on the page. You can edit these fields directly if you want precise positioning, or drag the preview image around on the page and watch the values change.

The *Center* buttons let you center the image *Vertically*, *Horizontally*, or *Both*. These buttons are particularly useful after adjusting *Size*, since you can re-center the image quickly.

## Making the Print: The Buttons at the Bottom

Once you're happy with the paper and resolution settings and the photo's size and position on the page, you can make a print using the *Print* button. You also have the option of saving the settings you just defined for *Size*, *Position*, and *Printer Settings* by clicking on *Save Settings*, or you can do both at once using *Print and Save Settings*. Any settings you save will be remembered for future GIMP sessions.

The other buttons—*Help*, *About*, and *Cancel*—are self-explanatory.

Once you click on *Print* or *Print and Save Settings*, GIMP should begin sending data to the printer. At a high-resolution setting, this may take a while; watch the status bar at the bottom-right of the GIMP image window for progress.

## Fine-Tuning: Image/Output Settings

Remember that *Image/Output Settings* tab that I skipped at the top of the Gimp-Print dialog? These settings let you fine-tune the brightness and color of the image, choosing from several presets. You can think of *Line art* as analogous to black-and-white drawings or text (2-bit color), *Solid colors* as analogous to indexed color (with areas over which the color is constant), and *Photograph* as 24-bit true color. If in doubt, leave at the default setting, *Photograph*.

*Output type* is more useful. If you're printing a grayscale or black-and-white image, selecting the appropriate value will restrict your printer to only black ink. This may be cheaper, depending on the printer; certainly it looks different from a monochrome image printed with colored inks. Try both and see the difference.

But the real power of the *Image/Output Settings* tab is the *Adjust Output...* button, which brings up the Print Color Adjust dialog shown in Figure 12-3.




Figure 12-3. The Print Color Adjust dialog

Here is where you adjust the brightness, contrast, and color balance of your prints.

Every printer model is a little different. Some may emphasize one color or another; others may print too light or too dark, or have other problems. This isn't very important when you're printing out web maps and directions, or even fliers for your company picnic. But when you're making glossy photos to hang on your wall or give to relatives, it matters.

To find out how your printer stacks up, you will probably have to experiment, printing several images and trying different settings. This can be slightly costly since you can use up ink and expensive paper. Fortunately, you don't need to print full-page images to test a printer: you can print small images, feed the same page through multiple times, and use the Gimp-Print position and size settings to place the image at a different location on the page each time. On a particularly tricky image, you might want to create a smaller test image that includes some of the "extremes" of the original.

The Gutenprint plug-in allows still more adjustments. It gives you a list of presets, such as *Mixed Text and Graphics* (the default) or *Line Art*, but selecting *Manual Control* gives you a long list of every value you might possibly want to adjust. Most of them are grayed-out by default; try clicking the toggle button next to any item to enable it.

Once you've found settings that look the way you want them (which usually means that the printed image looks similar to the image you see on screen), you can *Save Settings*, and Gimp-Print will remember these settings next time.

Each setting has a tooltip: to see it, hover over the text field for that setting. Some of the tricky ones are *Density*, *Gamma*, and *Dither Algorithm*.

*Density* controls the amount of ink used for printing. Normally, this should not need adjustment, but if your prints come out with so much ink that it smears or soaks into the paper, or so little that black areas aren't fully black or colors seem faded, this is the place to adjust it.

You met *Gamma* briefly in Chapter 2: it corresponds to the middle slider in the Levels dialog. Generally, gamma works like brightness: larger values mean a brighter print; smaller values mean a darker one. However, adjusting gamma won't change the levels of black or white in the image, only the middle tones in between.

*Dither Algorithm* controls how the GIMP maps its internal RGB colors to the CMYK colors the printer understands. In addition, printing almost always involves some scaling: it's not likely that the number of pixels in your image will exactly match the number of dots your printer would use at the size you want the image to appear.

Fortunately, the print software does a good job of scaling and color mapping. Most of the time, the default methods it uses are fine, but if you ever want to fine-tune the result, you can choose from a variety of algorithms: *Adaptive Hybrid* (the default), *Ordered*, *Fast*, *Very Fast*, *Hybrid Floyd-Steinberg*, or *EvenTone*.

The Gimp-Print project recommends using *Adaptive Hybrid* for highest quality. That should be the default, and most of the time you can leave it alone. If you want to experiment with other settings, *Ordered* can work just as well and is faster for images with mostly continuous tones (no text or line art). *Fast* is faster, but at the expense of quality (*Very Fast* is more so, but it sacrifices more quality than most people will tolerate). You can get details on some of the other algorithms on the Gimp-Print website, [gimp-print.sourceforge.net](http://gimp-print.sourceforge.net).

## Print to File

Even if you don't have a printer connected to your GIMP machine, you can print to a PostScript file. In the *Printer name* menu at the top of *Printer Settings*, choose *\*File*.

You will be prompted for a file name. It's probably best to give the file a name with a .ps extension, so that every machine where you try to use it will recognize that it's PostScript.

## Screen Shots and Scanners: the Acquire Menu

Most of your work with GIMP will probably begin with existing images, such as photographs, or a blank canvas.

But there's another way to obtain images: the *File* ► *Acquire* menu.

*Acquire* offers two seemingly unrelated functions: screen shots, and communication with devices such as scanners. But they both have something in common: GIMP needs to run a system program in order to import an image.

The *Acquire* menu also offers *Paste as New*, which you've probably already used in other contexts. It simply creates a new image out of whatever layer or selection has been most recently copied or cut.

### Scanning from GIMP

If you don't have a scanner, only *Screenshot* will be available in the menu. If you have a scanner, you need two software pieces in order to scan directly to the GIMP:

- The driver for your scanner
- A scanning plug-in for the GIMP

Scanner drivers use one of two common protocols: TWAIN or SANE.

TWAIN is most commonly used on Windows systems, and is also popular on Mac. It's usually provided as part of the drivers shipped with a new scanner. The origin of the name is unclear: some say that it stands for "Technology Without An Important Name," but people on the TWAIN committee say it comes from Kipling's phrase "Never the twain shall meet," reflecting the previous difficulty of connecting scanners to computers.

SANE is an open source protocol. It stands for "Scanner Access Now Easy." It's commonly used on Linux and Mac systems, but it's available for Windows as an add-on. Its home page is <http://www.sane-project.org/>.

Macs add one more menu entry, *Image Capture...*, which can grab images from a webcam, if you have one.

Whichever protocol you use for your scanner, you will need an appropriate plug-in for the GIMP to find the scanner. For SANE, there's the *xsane-gimp* plug-in; for TWAIN, there's a *Twain Acquire* plug-in that should be part of the GIMP packages on Mac and Windows.

Either of these plug-ins will create a new submenu inside the *Acquire* menu, called *XSane* or *TWAIN Acquire*. With your scanner turned on, check the contents of this menu. It may list the scanner there already, in which case the menu item will bring up a scanning dialog. Alternately, it may simply give you the option to bring up a device dialog, which will attempt to identify your scanner.

Scanners are tricky beasts. Some scanners work perfectly out of the box, but others require help. You may have to fiddle with your scanner setup to make it show up as a TWAIN or SANE device. Possible issues with specific scanners are beyond the scope of this book.

### Screen Shots

Screen shots are far less problematic than scanners, since there's no device driver to cause difficulties. *File* ► *Acquire* ► *Screenshot...* brings up the dialog shown in Figure 12-4.

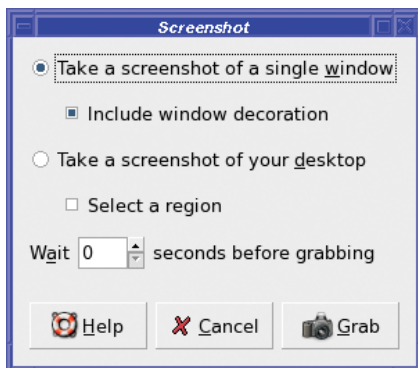


Figure 12-4. The Screenshot dialog

The options are all straightforward. You can take a shot of a single window, with or without its borders: click *Grab*, and then click on the window you want to shoot. *Include window decoration* controls whether the screen shot will include the window's frame and title bar. Or you can take a shot of your whole desktop. You can also choose *Select a region* in order to drag out a region of the desktop with the mouse.

The *Wait [0] seconds before grabbing* is particularly useful if you need to perform an action in the window while the screen shot is happening. For example, if you want a shot of a window with a menu showing, you can set a delay of five seconds, click *Grab*, move to the target window and click to select it; then you have five seconds to pop up a menu. You can set the delay as long or as short as you like, depending on the difficulty of the operation (and the speed of the system).

Since GIMP calls a system program to perform the actual screen shot, there are some limitations. For instance, on some systems, screen shots don't show the mouse cursor. If you need a screen shot showing a cursor, you'll have to cheat by getting a copy of the cursor image somewhere else and pasting it into your screen shot.

---

**Tip** If you ever need GIMP's special cursors for a screen shot, the GIMP source includes a file called `gimp-tool-cursors.xcf` which includes all the cursors used in GIMP. Each part of each cursor is a separate layer, so you can combine them as needed.

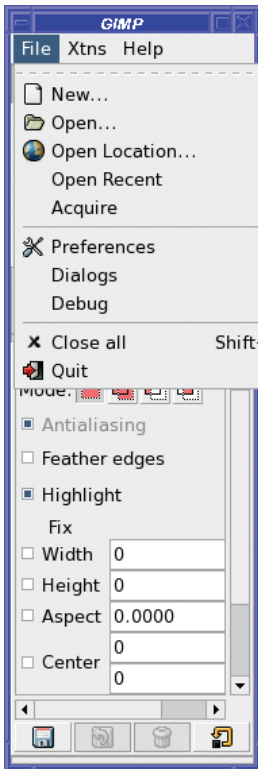
---

Screen shots of windows are usually limited to the window's own boundaries (Figure 12-5). If your menu goes beyond the boundaries of the window, you're better off shooting the whole desktop, or a region of the desktop (you'll have to guess how far the menus will extend). You can always crop off any extra.

Certain types of windows, such as games that are using accelerated 3-D graphics, may not produce a screen shot that looks exactly like what you see.

Finally, GIMP will shoot the content of windows as it appears on the screen. If your window is slightly off the screen, or is partially covered by another window, the screen shot will show that.





**Figure 12-5.** A screen shot of a window is limited to the window's boundaries, even if menus are showing that extend beyond the boundaries.

## GIMP Configuration Files

By now, you've probably referenced your GIMP profile directory for some reason or other—to install scripts or plug-ins, or to create new brushes or patterns. In addition, you've probably fiddled with your preference settings through GIMP's Preference dialog.

All these changes are saved inside your profile folder. But there's quite a bit more buried in there. Sometimes it's useful to examine your settings directly, or to change settings that aren't visible in the Preferences window. In addition, you can copy your profile folder from one machine to another if you want to make your GIMP settings available on a second machine.

Remember, you can find your GIMP profile folder by opening the Preferences window (*File* ► *Preferences* from the Toolbox window) and looking in the *Folders* category. Nearly every item in that category will have two entries: one for a system location, and one inside your profile folder. Take a look at your profile folder now as you go through its contents.

All of the configuration files in your GIMP profile folder are readable text files (though there are other types of files there as well, such as C plug-ins, brushes, patterns, and so on). You can copy your whole GIMP profile, or any part of it, to another machine (though compiled C plug-ins may not work if you move them to a different platform).

## gimprc

Most of your saved preferences are stored in a file called `gimprc`. If you've changed any preferences from the default, the file might include lines like these:

```
# GIMP gimprc
#
# This is your personal gimprc file. Any variable defined in
# this file takes precedence over the value defined in the
# system-wide gimprc: /etc/gimp/2.0/gimprc
# Most values can be set within The GIMP by changing some
# options in the Preferences dialog.

(interpolation-type lanczos)
(default-font "Helvetica Bold Italic")
(default-image
  (width 500)
  (height 400)
  (unit pixels)
  (xresolution 72.000000)
  (yresolution 72.000000)
  (resolution-unit inches)
  (image-type rgb)
  (fill-type background-fill)
  (comment "Created with The GIMP"))
(undo-levels 11)
(resize-windows-on-resize yes)
```

You'll probably recognize most of the options as preferences that are settable in the Preferences window. But what's that *default-font* line? Try as you might, you won't be able to find a default font setting in the Preferences window.

The GIMP can understand preferences beyond the ones exposed in the Preferences window, and editing `gimprc` is the way to set them. As with editing scripts or plug-ins, you should use a plain text editor to edit `gimprc` (or any other file inside your GIMP profile), not a word processor (Chapter 11, “Why a Text Editor?”, explains why).

First, you've probably already noticed that the syntax looks like Script-Fu: everything is delimited by parentheses. Options that take multiple values, such as the `default-image` settings, take a list of parenthesis-delimited lists.

You may wonder about that comment referencing `/etc/gimp/2.0/gimprc`: why is it 2.0, and not 2.2 or 2.4? GIMP sometimes confusingly uses “2.0” to refer to resources such as plug-ins or preferences that can be used by any GIMP version starting with 2.

Linux and Unix users can get the latest on `gimprc` by typing **man gimprc** into a terminal window (it also works on Macs if you used DarwinPorts to install the GIMP). This displays the manual page for `gimprc`, updated as of your current GIMP version, and should tell you all the latest options. On other platforms, or if you don't have the manual page installed, you can read it online at <http://www.gimp.org/man/gimprc.html>.

How do you find out which options you can set in `gimprc`? Here are some of the highlights that aren't already accessible through the Preferences window.

## Stingy Memory Use

Any large program like the GIMP needs to make trade-offs between speed and memory usage. By default, GIMP will assume there's plenty of memory, and will optimize for speed. On a low-memory system, you may see better performance by setting:

```
(stingy-memory-use yes)
```

## Default Brush, Pattern, Palette, Gradient, and Font

The GIMP doesn't normally remember settings for these values (though some of them will be remembered if you chose *Save tool options now* from the *Tool Options* category of the Preferences window). But you can set them explicitly in `gimprc`, with lines such as the following:

```
(default-brush "Circle (4)")  
(default-pattern "Leopard")  
(default-palette "Web")  
(default-gradient "Full Saturation Spectrum CCW")  
(default-font "Helvetica Bold Italic")
```

The names used in `gimprc` correspond to the names you see in the dialogs where you choose these items.

## Share Palette

The *Tool Options* category of Preferences lets you share brushes, patterns, or gradients, among all tools. But the GIMP can also share fonts and palettes:

```
(global-font yes)  
(global-palette yes)
```

## “Perfect Mouse”

When you move the mouse quickly, computer systems may take a shortcut and pass only some of the positions to the application. This saves GIMP from needing to process every pixel that the mouse covers, and usually speeds up drawing.

Normally, GIMP doesn't take this shortcut: it processes every pixel the mouse crosses. If you notice lags while drawing lines quickly, it might be worth experimenting with the *perfect-mouse* setting by changing it to “no”:

```
(perfect-mouse no)
```

Imperfection may be faster on some systems.

## Default Threshold

Tools such as Fuzzy Select and Bucket Fill have a threshold for what constitutes a “nearby” region, settable in Tool Options. Normally, this defaults to 15, but you can change that:

```
(default-threshold 10)
```

## Saving Modified Images

Usually, when you tell GIMP to save, it will save even if it thinks the image hasn't been modified. You can prevent this behavior if you like:

```
(trust-dirty-flag yes)
```

## Help Browser on F1

You can disable the Help browser (for instance, if you don't have one installed and want to use F1 for something else):

```
(use-help no)
```

## Help Locales

Normally, GIMP should be able to figure out which language you use from your system's "locale" setting. If this doesn't work, though, you can specify it with `help-locales`. This is a colon-separated list. For instance, if you have US English plus a Great Britain locale that uses the UTF-8 character set, use the following:

```
(help-locales "en_US:en_GB.UTF-8")
```

## Restoring Keyboard Shortcuts

The GIMP has a preference in the *Interface* category to *Save keyboard shortcuts on exit* (assuming you've enabled dynamic keyboard shortcuts in the first place, of course). Naturally, GIMP will usually load your saved shortcuts whenever it starts. But if you don't want that to happen automatically, add this to your `gimprc`:

```
(restore-accel no)
```

In that case, GIMP will only load your customized key bindings when you go to the *Interface* category of the Preferences window and click on *Reset Saved Keyboard Shortcuts to Default Values*.

## Tear-Off Menus

If for some reason you don't like the tear-off menus, you can disable them:

```
(tearoff-menus no)
```

## Plug-In Paths

You saw the Fractal Explorer, Gfig, Gflare, and GIMPressionist plug-ins in Chapter 7. You can set their data directories to something other than the default, or add additional directories:

```
(fractalexplorer-path
  "${gimp_dir}/fractalexplorer:${gimp_data_dir}/fractalexplorer")
(gfig-path "${gimp_dir}/gfig:${gimp_data_dir}/gfig")
(gflare-path "${gimp_dir}/gflare:${gimp_data_dir}/gflare")
(gimpressionist-path
  "${gimp_dir}/gimpressionist:${gimp_data_dir}/gimpressionist")
```

You can include multiple folders by using a colon between the names. `${gimp_dir}` refers to your personal `gimprc` directory, and `${gimp_data_dir}` to the GIMP system directory where these plug-ins are stored.

## Changing Preview Sizes with `gtkrc`

You may remember from installing GIMP that it's built upon a user interface library called GTK, the Gimp Toolkit.

GTK's numerous configuration options can be set globally, but you can also set options for specific applications that take precedence over the global settings. GIMP's GTK settings reside in a file called `gtkrc` in your GIMP profile folder.

Most of the time, this will be updated automatically to reflect changes you make in your GIMP user interface. But there are a few cases where it's helpful to view the file.

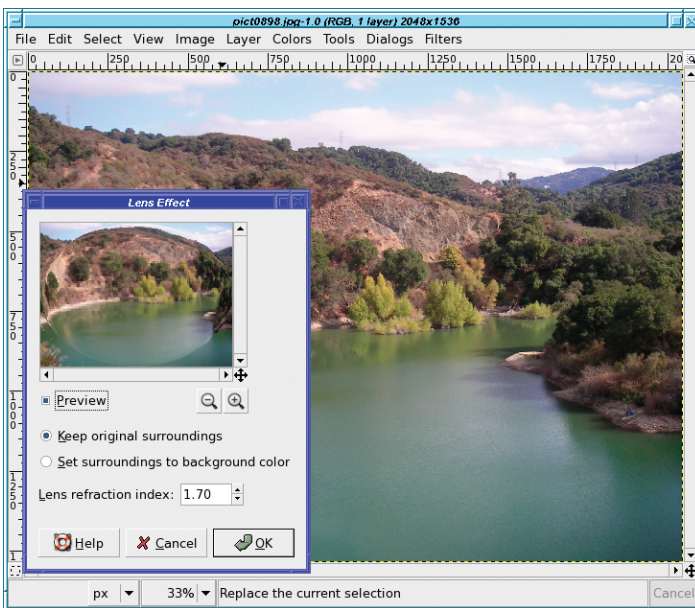
The most important use of `gtkrc` is to change the preview size in plug-in dialogs.

You may have puzzled over the two *Preview Size* options in the *Interfaces* category of Preferences. When you change the size, the previews in plug-in windows don't change. Why not?

The two options in Preferences affect different previews. *Default layer & channel preview size* controls the size of the preview thumbnails shown in each line of the Layers dialog. But they're tricky to use: GIMP doesn't actually honor them unless you close the tab, bring up the dialog again, and then re-dock it.

It's much faster to ignore the preferences and change the preview sizes from the dockable's menu. (You can also change them by editing `sessionrc`, as you'll see shortly.) *Navigation preview size* controls the size of the window that pops up when you click on the navigation panner button at the bottom-right of an image window.

So what about those plug-in previews? The normal preview size is pretty small and it's sometimes hard to see much (Figure 12-6). Can you make it bigger?



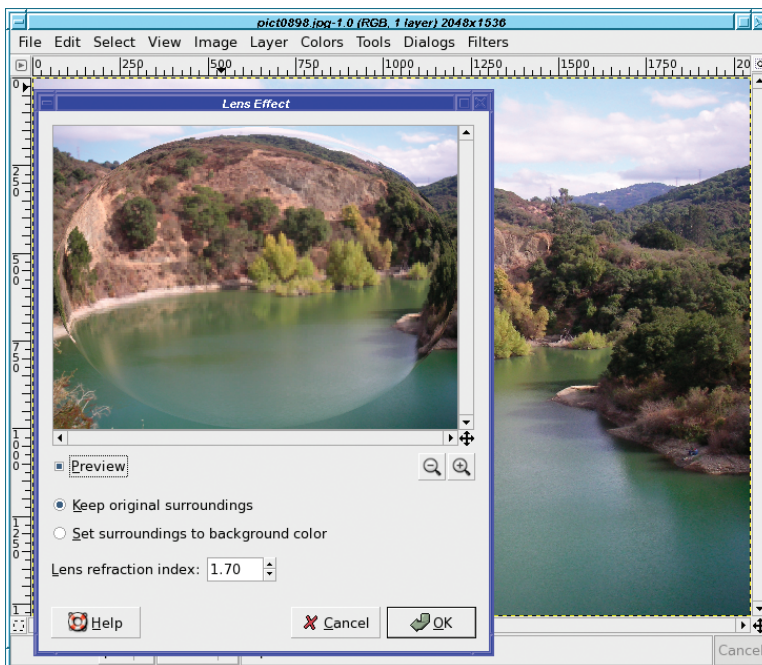
**Figure 12-6.** GIMP's default preview size is a bit small.

Not through the Preferences window, but you can if you're willing to edit `gtkrc`. (If you don't already have a `gtkrc` file in your GIMP profile folder, you can create a new one with your text editor.)

Add some lines like this:

```
style "gimp-large-preview"
{
  GimpPreview::size = 350
}
class "GimpPreview" style "gimp-large-preview"
```

In place of 350, put whatever size you want your previews to be. The default is 200. You don't have to use the name `gimp-large-preview`; any name will do. As with most configuration file changes, you'll have to restart GIMP to see this change (Figure 12-7).



**Figure 12-7.** Increase the preview size as much as you like.

## GIMP's Window Positions and Configuration: `sessionrc`

The GIMP remembers the positions of your windows, so you can stow the Toolbox on one side of the screen and the Layers dialog on the other. It stores this information in the `sessionrc` file.

You will seldom need to touch this file. But this is the place to look if you have a problem with your window positions, or with the preview sizes in dockables like Layers. You can edit these values directly in `sessionrc`. Don't edit this file while GIMP is running, though, or GIMP will probably overwrite your changes.



## Menu Entries and Key Bindings: `menurc`

The `menurc` file is the place where GIMP stores any key bindings you assign to menu items.

GIMP fills the file with all of its menu items—many hundreds of them—but most of them are commented out with a semicolon at the beginning of the line. A line that's commented out means that it hasn't changed from GIMP's default setting. Anything that's not commented is a key binding that you've added.

Key bindings look like this:

```
; (gtk_accel_path "<Actions>/file/file-save" "<Control>s")
; (gtk_accel_path "<Actions>/file/file-save-as" "<Shift><Control>s")
(gtk_accel_path "<Actions>/plug-in/plug_in_screenshot" "<Alt>s")
(gtk_accel_path "<Actions>/view/view-zoom-in" "equal")
; (gtk_accel_path "<Actions>/select/select-none" "<Shift><Control>a")
```

These lines say that Alt+S has been added as the key binding to do a screen shot, and the = (Equal) key will zoom in. Save, Save As, and Select None are at their default settings and haven't been changed.

Unfortunately, you can't use this file to set two different bindings for the same function (for instance, to make either = or + zoom in). Only the last binding you set will be used.

You shouldn't normally need to edit `menurc`. But it's a useful place to look if you're ever curious about which key bindings you've changed. The GIMP's Configure Keyboard Shortcuts dialog (accessed from Preferences) will let you change or remove individual bindings, but it doesn't offer a way to find out what you have changed from the default settings.

There's one other use for `menurc`: you can quickly install a whole collection of bindings. For instance, the GIMP comes with a file called `ps-menurc` that makes GIMP's key bindings more like Photoshop's. Although it's usually a better idea to learn how to use GIMP with its native settings, people who are used to Photoshop may find it easier to learn GIMP starting with a familiar set of bindings. To use a file like this, quit GIMP, copy `ps-menurc` as `menurc` in your GIMP profile and restart GIMP, and you'll have the new bindings. You can also make your own custom key-binding set, and distribute them to other people.

Of course, if you've already made changes to your GIMP key bindings, it's a good idea to back up your `menurc` somewhere before you replace it with another file, just in case you want to go back. If you ever want to return to GIMP's default setup, just remove any `menurc` file in your profile. GIMP will regenerate it with default values (everything commented out) the next time you start up.

## Other Files

- `colorrc` is where GIMP stores the recently used colors that show up in the color chooser window.
- `controllerc` has information about your mouse and keyboard—the items described in the *Input Controllers* category of Preferences.
- `devicerc` holds the information saved when you click *Save Input Device Settings Now* in the Preferences window's *Input Devices* category. It includes information about your graphics tablet if you have one; but in addition, it saves tool settings for the drawing tools.

- `documents` is the list of images you've edited recently: the same list you see in the *Open Recent* menu.
- `gimp-splash.png` is a way to set a custom “splash image” for GIMP to use when it starts up. This file isn't created automatically, but any PNG image you put here will become GIMP's new splash image. (See also the `splashes` folder, discussed below.)
- `gtk-bookmarks.sav` holds the bookmarks you set in the *Open* and *Save As* dialogs.
- `parasiterc` holds global *parasites*. A parasite is a name GIMP gives to any sort of persistent data. Plug-ins can use this file to store default values.
- `pluginrc` is a list of all your plug-ins and the parameters they take. Editing this file is almost always a bad idea, and GIMP regenerates it frequently (whenever it finds a new plug-in). The actual plug-ins are stored in the folder called `plug-ins`.
- `printrc` is where the Gimp-Print plug-in stores your printer settings.
- `templaterc` stores information about any templates you have, and GIMP's standard ones. The actual templates, if you've created any, are stored in the `templates` folder.
- `themerc` stores information about the theme you're using, while themes you import should be stored in the `themes` folder.
- If you've customized your Toolbox, that information will be stored in `toolrc`. Even if you haven't, the standard Toolbox setup is stored in that file.
- `unitrc` contains a list of the units GIMP knows about. You shouldn't need to edit it.

## Folders Inside the GIMP Profile

- Brushes, curves, gradients, levels, palettes, patterns, templates, and themes all offer places to store tools you make in GIMP.
- `Fractalexplorer`, `gfig`, `gflare`, and `gimpressionist` are used by these four plug-ins to store their options and data files.
- In `environ`, you can create files setting variables that you want plug-ins to be able to use. For instance, Python needs to know where the actual Python binary is on the system, which it finds in a variable called `PYTHONPATH`. So if Python is installed, you should have a file in the system-wide `environ` folder with a line like:

```
: PYTHONPATH=/usr/lib/gimp/2.0/python
```

The colon tells GIMP to prepend the value to any existing `PYTHONPATH` that might already be defined on the system, using the colon as the separator character.

- The `fonts` folder is a place where you can install fonts that you only want to use in GIMP, not in other programs.
- `interpreters` is a place where you can store command interpreters other than Script-Fu. This is mostly for use in scripting. See your system-wide `interpreters` folder for examples of how these files look. It's not likely most people would try to add one: you'll only want to put files here if you're creating GIMP bindings in a new language.

- GIMP modules are loadable executables, somewhat like plug-ins, but at a lower level. The `gimp-perl` and `gimp-python` interpreters, which make it possible to write plug-ins in those languages, are modules. If you developed your own GIMP module, this is where you'd put it.
- The plug-ins and scripts folders are where you should install plug-ins (C, Perl, or Python) or scripts (Script-Fu).
- The splashes folder isn't created automatically. But if you create a folder by this name and put images in it, when GIMP starts up it will use one of your images (chosen at random) as its "splash" screen, rather than the built-in splash image. The images can be in any format except XCF.
- The tool-options folder is where the tool options for all of GIMP's tools are stored.
- `tmp` is simply a folder GIMP uses when it needs to create any sort of temporary file. If GIMP ever crashes, you may end up with some extra files here.

## Additional Resources

You're nearly at the end of the book! Where do you go from here?

Fortunately, there's an extensive collection of GIMP information on the web, on the GIMP website, and elsewhere. If you want to learn more about GIMP, you won't run out of reading material any time soon!

## Official GIMP Documentation

The GIMP is one of the better-documented free software projects. To begin with, it has an extensive reference manual online: the GIMP User Manual. Point your browser at [gimp.org](http://gimp.org) and click on "Documentation" to see the manual. It's available in English, Chinese, Czech, Dutch, French, German, Italian, and Swedish.

Whenever you're unsure about how to use a particular feature or what an option means, and you don't find the answer in this book, try the GIMP User Manual. It has a few gaps (remember, it's all written by volunteers, primarily the GIMP developers themselves), but it's amazingly thorough.

The manual begins with a "What's New?" section, which is a useful reference for finding out what has changed from one major version of GIMP (2.0, 2.2, 2.4) to the next.

### WHY DO GIMP VERSIONS CHANGE BY .2 INSTEAD OF .1?

The GIMP, like many free software projects, uses the convention that an odd version number after the decimal point is used for a *development version*, while an even number means a *stable version*.

After 2.2 was officially released, the developers began work on 2.3. That version number has been used for the unstable development version as new features are stabilized and important bugs fixed. When it's stable, the version number will be changed to 2.4, to indicate a new stable release.

While 2.3 was being developed, any important bugs found in 2.2 were fixed, resulting in new GIMP stable versions such as 2.2.1, 2.2.2, and so on. Most users used one of the stable 2.2 releases, while people who wanted to help out with GIMP development could download the latest 2.3 and try it out, reporting bugs if they found any.

Development releases are usually source releases. In order to try the very latest features or help with finding bugs, you may have to compile the GIMP yourself.

## Other Official GIMP User Documentation

Beyond the GIMP User Manual, the GIMP website has quite a collection of other documentation. This is the best place to find documents on GIMP configuration (how to install new fonts, tips for choosing your tile cache size to get the best performance, and so on).

The GIMP's manual pages are also available online in the “Documentation” section of the GIMP website. If you're not on a Unix-like machine and don't have the manual pages installed, but you want to read about the options you can set in `gimprc`, or about the details of building and installing plug-ins with `gimptool`, you can find all the information here. There's also documentation on the Script-Fu and Python modules and some information on the Gimp-Print plug-in.

But perhaps the most useful are the tutorials. The official *gimp.org* tutorials are divided into categories: mostly by experience level (Beginner, Intermediate, Expert), but also by category (Photo Editing, Web, Script Authoring).

Want to know how to make text that looks like it's made of shining gold? To find ways of making your photos look like sketches, beyond the simple layer-mode technique of Chapter 10, or to turn a sketch into a colored cartoon-style figure? Want to see how a real artist uses GIMP techniques to draw a paintbrush with shading and a polished wood handle? Want to make a movie from a selection of images, or eliminate the noise from your camera's CCD chip? These techniques, and a lot more, are all on the GIMP tutorial site on pages that take you through the process, step by step.

## Developer Documentation

Finally, there's documentation for developers at *developer.gimp.org*, which includes tutorials on writing scripts, articles about writing C plug-ins, and a fairly complete and very readable set of documentation on the GIMP's API (*Application Program Interface*, the set of routines called by GIMP plug-ins and internal functions to do their work). This website is an indispensable resource for anyone writing GIMP plug-ins or contributing code to the GIMP.

Equally important is the *Developer FAQ*, which lists guidelines for people interested in contributing code to the GIMP. Potential developers should also read the link labeled “Hacking,” which gives you information on which packages you're likely to need in order to build the development releases and about some of the coding standards. (This is just a copy of the file `HACKING` which comes as part of the GIMP source.) And there's information about how to check out the very latest version of GIMP—the one the developers are working on *today*—from CVS (the *Concurrent Version System*).

A link labeled *Standards* gives an excellent discussion of the various image-format standards GIMP supports, as well as other standards covering such topics as color management, desktop icons, thumbnail management, and how drag-and-drop is implemented.

One of the fun entries on *developer.gimp.org* is the *ChangeLog*. This is a file in the GIMP source where every change made to the GIMP is recorded. If you're curious about the pace of GIMP development, or about who works on which tasks, the *ChangeLog* can be interesting reading. Of course, the comments in it are usually deep in technical detail: you won't find a lot of high-level discussion of new features or plans.

*News* is a higher-level version of the same thing. Instead of showing every check-in by every developer, it summarizes the changes from one version to the next, much like the “What's New” section of the User Manual, except that it covers the development *point releases* such as 2.3.3 and 2.3.4. Even if you're not a developer and not building GIMP from source, it's sometimes enlightening to check the developers' news and see what features you can look forward to in the next major version.

## The Wiki

In addition to the official documentation, GIMP also has a Wiki, located at *wiki.gimp.org*.

A *Wiki* (the name comes from the Hawaiian term *wiki-wiki*, meaning *quick*) is a web application that allows users to contribute their own pages—or add to existing pages—to build a knowledge base. Wikipedia is the best-known example, but many projects like GIMP have their own Wiki to let users contribute new documentation and hints.

The GIMP Wiki currently contains a good list of tutorials (including many links to tutorials on other sites), some user FAQ (Frequently Asked Questions) documents, hints for printing and scanning, hints on building GIMP and on reading its internal source code, notes on navigating the GIMP website, and many more topics. It's a good place to poke around.

## Other Tutorial Sites

Of course, the GIMP information available on the web isn't limited to the official documentation on the GIMP website. GIMP is used by countless people worldwide, and many GIMP users enjoy sharing tips they've discovered or techniques they've learned.

You've probably already looked at the tutorials on *gimp.org*. But there are hundreds of tutorials on the web discussing how to use GIMP to accomplish various tasks.

If you're looking for a really specific effect, try a web search. Suppose you want to add fire to a drawing you're making. Try searching for **gimp fire** or **gimp fire tutorial** and see what comes up. Often you'll get several different hits, so you'll end up with a choice of tutorials showing different ways of accomplishing the same effect. Choose the one you like best, or combine them.

You can usually adapt tutorials that were written for other programs, too. If you search for GIMP tutorials and don't find any, try searching for something like **fire image tutorial** or even **fire photoshop** or **fire "paintshop pro"**. There are lots of web tutorials written for other programs, and many of them work just fine in GIMP as long as you don't expect functions to be in the same places in the menus, use exactly the same settings, or have specific key bindings.

If you don't have a particular effect in mind and just want to learn something new or get ideas, try browsing tutorials. Once you've exhausted the ones on the GIMP site, there are many more sites for finding GIMP lessons.

Of course, the nature of the web being what it is, some of the tutorials will be old and written for older versions of the GIMP, such as GIMP 1.2. Don't let that scare you off, though. The techniques should all work in any version, so at worst, you might have to spend a little time searching through the menus or in the Plug-in Browser window to find a function mentioned in a tutorial.

Some good sources for GIMP tutorials are as follows:

The GIMP Wiki's Tutorial page:

<http://wiki.gimp.org/gimp/GimpTutorials>

Tigert Labs ("Tigert," you may remember from Chapter 1, is Tuomas Kuosmanen, GIMP contributor and the artist who created the GIMP mascot Wilber):

<http://tigert.gimp.org/gimp/tutorials/>

The GIMP User's Group, or GUG (which also has a discussion mailing list and a place for artists to upload their images):

<http://gug.sunsite.dk/?page=tutorials>

The WinGIMP site (written for Windows users, but the tutorials apply to every platform):

<http://www.wingimp.org/tutorial/>

GimpTalk:

<http://www.gimptalk.com/>

This site has a collection of tutorials from all over:

<http://empyrean.lib.ndsu.nodak.edu/~nem/gimp/tuts/>

In addition, there are some videos illustrating GIMP features at:

<http://jimmac.musichall.cz/gimp2demos.php>

Don't stop there, though. If you don't like any of those sites, try a web search for **gimp tutorials** or **gimp tutorial** and see how many other interesting sites come up. You may run out of time to read tutorials, but you probably won't run out of tutorials to read.

## Mailing Lists

If you want to talk with other GIMP users and exchange tips, you can choose from quite a selection of mailing lists. Just go to [gimp.org](http://gimp.org) and click on "Mailing lists."

Your choices include the following:

- *GIMP User*: Discussions and questions related to using the GIMP.
- *GIMP Announce*: Announcements of new releases and other important GIMP news, without any discussion.
- *GIMPwin Users*: Discussions for GIMP users on Windows. A great place to go to if you need help with Windows-specific issues such as printing or building plug-ins.
- *GIMP Developer*: Discussions of the development of the next GIMP version.
- *GIMP Web*: Maintenance of the GIMP website.
- *GIMP Docs*: Maintenance and development of GIMP documentation, especially the User Manual.
- *GIMP Images*: A place for people to exchange images created with GIMP.
- *GIMP Perl*: Use of the Perl module to develop plug-ins.
- *GIMP OS2*: Discussions related to using GIMP on IBM's OS/2 operating system.

The mailing list page includes tips on "netiquette." If you haven't spent a lot of time on technical mailing lists, it's a good idea to read the suggestions. It's often helpful to "lurk" for a while on a list before making your first posting, to become familiar with the culture of the list.

In addition to the official lists on the GIMP mailing list page, the GIMP User Group has several online forums, plus a mailing list for artists using GIMP. Visit <http://gug.sunsite.dk/> and click on "Mailing list."



## Reporting Bugs

What if you're pretty sure you've found a bug in GIMP?

The GIMP's bug system is entirely open. Anyone can report a bug, or search for known bugs. The project uses a bug-tracking program known as Bugzilla, located at <http://bugzilla.gimp.org>.

When you first go to GIMP's Bugzilla page, you will be presented with a login screen and a chance to create an account. You can search the bug system all you want without creating an account, but to file a new bug or add comments to an existing one, you'll need to register. That's so that the Bugzilla system can notify you by email if the bug is fixed, if a developer needs more information from you, or if anyone else adds a comment regarding your bug.

It's a good idea to search Bugzilla before reporting a new bug. It may be that someone else has already discovered your bug and reported it. To do that, click on Search. *Don't Panic!* The Search page looks scary, but if you ignore most of the page and use only a couple of the fields, it's actually quite simple.

First, you need some keywords for the bug you're looking for, just like you would for a web search. If you think you've found a bug in the Zealous Crop plug-in, then searching for "zealous crop" is a good start (you don't need to worry about capitalization—Bugzilla searches are case-insensitive).

Then go to the Product field on the left side of the window and scroll down to find GIMP (Figure 12-8).

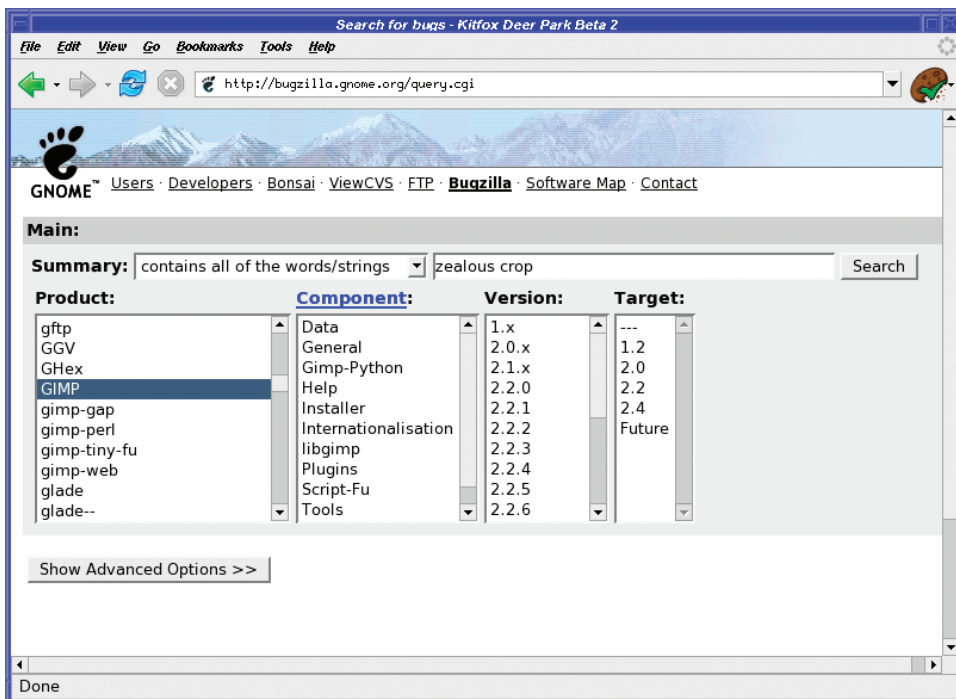


Figure 12-8. The Bugzilla search page

You may be amazed by the long Products list. What are all those other products?

GIMP shares its Bugzilla system with a project called Gnome, which includes GTK (the toolkit GIMP uses) as well as many other applications. The total number of programs covered by the Gnome bug database is huge, which is why it's important to select the GIMP "Product" and narrow the search to only GIMP bugs.

Once you've typed some keywords in the *Summary* and selected the GIMP project, you can click the *Search* button (to the right of the *Summary*; depending on your window size you may need to scroll right to see it). Or just hit the Enter key. You'll be taken to a page showing a list of GIMP bugs related to your keywords. If there are no bugs associated with those keywords, Bugzilla will say "Zarro Boogs found." That was an overworked programmer's idea of a joke early in the development of Bugzilla, and unfortunately it stuck.

You can always use the *Back* button in your browser to return to the search page and try a different set of keywords. If you're surprised to find "Zarro Boogs," it might be worth clicking on the *Show Advanced Options* button and doing another search with all the *Status* options selected (Shift+click and Control+click in these lists will let you highlight multiple values). By default, Bugzilla searches only for bugs that are still "Open." When a bug has been fixed in the current development version of GIMP, the bug is marked "Resolved" in Bugzilla, even though the version you're using may still have the bug.

If you find an existing bug that's similar to yours, read it. If you want to be notified when the bug is fixed or when someone makes a comment, add your email address (the one you registered with Bugzilla) in the *Add CC* field. If you have more details, you can add them: Scroll down to *Additional Comments* and type in your information. When you're finished, click *Commit*.

Reporting a new bug is slightly more complicated. When you go to [bugzilla.gimp.org](http://bugzilla.gimp.org), you're presented with the bug-reporting form. Otherwise, click *New*.

The most difficult part of bug reporting may be choosing the right *Component*. GIMP has quite a few components, such as *Data*, *General*, *Help*, *Installer*, *Internationalization*, *Plug-ins*, and *User Interface*. If one of these categories seems relevant to your bug, choose it. If you're not sure, it's safe to choose *General*; someone who knows more about the structure of GIMP can change that later.

Be sure to specify the version of GIMP you're using. You can get that from *Help* ► *About*, or by typing **`gimp --version`** (that's two minus signs before "version") at the command prompt. Don't forget to mention which platform (Windows, Linux, and so on) you're using. You can use *Distribution/Version* to clarify that: for instance, if you choose *Windows* for Operating System, you might type something like **XP** or **2000** in the *Distribution/Version* field.

You can usually leave the *Severity* field set to *Normal*, unless you've found a way to make GIMP crash, in which case you can use *Critical*. Don't overstate the severity of your bug to try to get the attention of developers.

Try to be very specific about what you did, what you expected to see, and why it's a problem. Vague bugs saying things like "Crop is broken" will likely be closed or ignored. It's much better to create a very specific bug like "The Zealous Crop plug-in ignores layer offsets," with a sample XCF image attached and a full description of what the plug-in does versus what you think it should do. If an error message was involved, write down the exact text of the error message; don't just paraphrase it. If you can't reproduce the bug in order to see the error message again—if it's something that only happened once and never again—the bug may not be worth reporting. Sometimes running GIMP from the command line will give you additional information about what's wrong, which you can then paste into a bug report.

Gathering this information can take time, but think about it this way: if you're hoping the developers will volunteer some of their free time to fix the bug, it can help a lot if you spend some of your own time making their jobs a little easier.

If the problem only happens with certain images, it's very helpful if you can attach an image that the developers can use to reproduce the problem. (If they can't make it happen themselves, they won't be able to fix the bug.) If you can reduce your sample image to something relatively small and simple, that's ideal. To attach an image (or anything else besides text) to a GIMP bug, click on *Create a New Attachment*. If you're filing a new bug, you'll have to submit the basics of the bug first, and then go back to it and add the attachment.

Bugzilla can be intimidating the first few times you use it. But don't let it scare you. If you've found a GIMP bug, the developers want to hear about it so they can fix it.

## Finding Source Code at cvs.gnome.org

The GIMP's source code is stored using the “Concurrent Version System” program, or CVS, on a machine called cvs.gnome.org.

If you have CVS installed, you can check out the latest version of the GIMP—the one the developers are working on—at any time. The GIMP website has instructions on how to do that.

But even if you don't have CVS, and don't want to download the entire GIMP source code (about 15MB), you can view specific files and their revision histories using the web interface to CVS: <http://cvs.gnome.org/viewcvs/gimp/>.

You can click on folder names to navigate to those folders, or on file names to view the revision history of the file (all the changes that have been made to it recently). If you want to see the current contents of a file, you can click on *View* to get a pretty colorized version, or *Download* to get a raw copy of the file that you can save to your local disk. This is an easy way to get copies of plug-in source to use as a basis for your own plug-ins.

## Where to Find Freely Available Images

Sometimes you need a photo for a project, and you just can't find anything quite right in your own collection. Darn it, I wish I had my own photograph of an erupting volcano, or of the Sphinx—but my travel budget doesn't allow for either one.

What can a poor GIMP user do? Look on the web, of course!

The problem? Most images are copyrighted by the person who created them. It's not ethical—or legal—to use someone else's photographs in your own project without permission. If it's just for your own use, or for GIMP practice, you might be able to get away with it. But if it's an image you're planning to post on the web, distribute to other people, or especially if you're using it for any sort of money-making purpose, it's important to use images with copyrights that allow re-use.

Fortunately, there's a movement to make more forms of information freely available, and several projects have been created to help people share their images. In some cases, the images are free for non-commercial use, but if you intend to make money with the image, you must contact the person who owns the copyright. Other images are free even for commercial use.

In addition, quite a few US government agencies put their images in the “public domain,” meaning they are free to use for any purpose because the images have already been paid for with taxpayer money. Government sites can offer excellent images on quite a range of subjects, and you may not need to look any further.

The Creative Commons is a nonprofit organization launched by open source advocate Lawrence Lessig to enable the sharing of artistic works of all sorts, including images. You can read about the Creative Commons idea on their website, <http://creativecommons.org/>.

The website lets you search for images you can use in your own projects. Some images have restrictions; others are completely free for any use. You can specify what sort of license you need when you search for images.

If you use images from Creative Commons (or even if you don't), consider making your resulting art available for other people to use. The more people who contribute to the commons, the better it works for everyone.

Burning Well is a repository for public-domain images: images that can be used without any restriction. The website is <http://www.burningwell.org/>. You can search for images by keyword, or browse images by category. There are over a thousand images, and some of them are very good. They may have exactly what you need.

Stock.XCHNG is another free photo site. They have a large searchable collection of images organized by category, as well as forums for discussing photography at all levels of experience. The website is <http://www.sxc.hu/>.

## U.S. Government Sites

The US government is an excellent source of images you can use without copyright worries. It is considered good manners to give credit where it's due, and mention the folks who did the work. See the specific copyright information on each site for more details.

The National Aeronautics and Space Administration (NASA) has an excellent range of images of planets and other astronomical objects, as well as images of spacecraft. Some specific NASA image collections are found at the NSSDC Photo Gallery (a sort of "NASA's Greatest Hits" collection) at [http://nssdc.gsfc.nasa.gov/photo\\_gallery](http://nssdc.gsfc.nasa.gov/photo_gallery); and the Jet Propulsion Laboratories (JPL) Planetary Photojournal at <http://photojournal.jpl.nasa.gov/index.html>.

And don't miss the NASA Image Of The Day Gallery at <http://www.nasa.gov/multimedia/imagegallery/>. Or search for images by keyword at NIX: <http://nix.nasa.gov/>. Finally, NASA has a website devoted to 70mm Hasselblad camera images taken by astronauts on lunar missions at <http://www.lpi.usra.edu/resources/apollo/catalog/70mm/>.

The US Fish and Wildlife Service offers photos of wildlife and related scenery at <http://images.fws.gov/>.

The US National Oceanic and Atmospheric Administration, NOAA, has lots of good images at <http://www.photolib.noaa.gov/>. In addition to photographs of storms, coastlines, ships, and corals, they have collections for marine animals, space (atmospheric satellites), and other interesting topics.

The National Park Service has huge collections of stunning scenic shots, as well as wildlife and cultural events related to the parks. The website is currently at <http://photo.itc.nps.gov/storage/images/index.html>. If that changes, start at <http://nps.gov> and poke around. You'll find it.

## Art Sites and Imaging Contests

You're armed to the teeth with GIMP techniques, and itching for an excuse to use them. Even more, you'd like an audience to appreciate some of the funny images you're making. Where can you look for a community of like-minded people?

The GIMP User Group, GUG, has already been mentioned. It's at <http://gug.sunsite.dk/> and it's full of enthusiastic GIMP artists. Go to "Gallery" to see some of their work. If you have images to share, you can create your own gallery.

Deviant Art is a community of digital artists. It's not all GIMP users, but a number of GIMP artists hang out there. There's an enormous sampling of digital art to browse, discussion forums, and it's even possible to sell art. The website can be found at <http://www.deviantart.com/>.

Then there are the "Photoshop contests." Don't let the name put you off: you can use GIMP or any other program to enter. Photoshop contests are websites where participants create digital images according to a set of guidelines, and then everyone votes for the winners. Most often, the images are very entertaining. Even if you're not interested in entering a contest, it can be addictive to browse Photoshop contest sites. And entering the contests can be a great incentive to practice your GIMP skills.

Some sites use a theme, like "Extreme sports," and participants are free to create any image that illustrates the theme, such as someone juggling live tigers, or waterskiing on the back of a great white shark.

Other sites provide an image, and the goal is to modify the image in some (usually humorous) way.

Some contest sites are free to enter while others charge a small fee; some give prizes to the winners, others don't.

A web search for "photoshop contest" should turn up several likely candidates.

## Summary and Conclusion

That's it! You're on your own now. But you should have enough information to accomplish anything you want in the GIMP.

You can read and understand your GIMP configuration files, and edit them when you need to make changes. You can print your GIMP creations (assuming you have a suitable printer) and make screen shots of your GIMP windows for writing your own tutorials. If you find a bug, you can track down the details and report it.

If the skills you need for making an image aren't in this book, maybe they're covered in the GIMP User Manual on the GIMP website. Or maybe there's a tutorial to show you exactly what you need, either one of the official ones on [gimp.org](http://gimp.org), or somewhere else on the web.

Finally, you know some sources to look for raw images for your creations, images with appropriate licenses so that you can distribute your derived art. You've explored some art sites that can give you inspiration, and you have places you can post your art when you have something to show off.

The rest is just practice. The more you work with the GIMP, the easier it gets. Tasks that seemed impossibly confusing the first time you tried them (like making a selection with Paths) become second nature. You may have stayed away from plug-ins because you're not a programmer—but eventually you may change your mind and decide to try writing or modifying one. You may have joined a mailing list or two in order to ask for help . . . and before long, you're the one helping other people. You may even decide some day to become a GIMP contributor, writing code or documentation or web pages.

Proceed at your own pace, and work with the kinds of images that interest you most.

The only rule is: Be sure to have fun!







# Getting and Installing the GIMP

**Y**ou've decided to learn to edit images with the GIMP! But if it's not already on your computer, first you have to install it.

You may have Windows, Linux, a Mac, or another UNIX system such as FreeBSD. All of them can run the GIMP, but each has a different installation method.

First, read through this introduction to installation, and then proceed to the appropriate appendix for your system.

But whatever you do, be sure to read through Appendix F (“First Startup”) at least once. Starting the GIMP for the first time can be a little daunting. It asks you a slew of questions for which you cannot possibly have reasonable answers. “First Startup” is brief and to the point. You'll be glad you read it.

In *this* appendix, you'll learn about the following:

- Requirements
- Where to find the GIMP

The next three appendixes are specific to various operating system types.

- Appendix B: Installing on Windows
- Appendix C: Installing on Mac
- Appendix D: Installing on Linux or Unix

Since the GIMP is open source, you can compile a version tailored to your system.

- Appendix E: Building from Source

And once your installation is finished, be absolutely certain to check out the final appendix:

- Appendix F: First Startup

## Requirements

The GIMP is highly portable. It runs on most versions of Windows, Mac OS X, Linux, and quite a few versions of Unix. Whatever machine you're using can probably run the GIMP.

It doesn't require an especially fast machine. However, if you want to work with large images, you'll find that adding more memory will help a lot. A faster processor won't hurt, of course.

GIMP is a GTK+ application. That means its user interface—menus, buttons, and so on—is made using the “GIMP ToolKit.” GTK+ was originally developed in order to build the GIMP, but has since become well known on its own as the basis for hundreds of other cross-platform open source applications. If you don’t already have GTK+, you’ll need to install it: the installation instructions for each platform will discuss where to get it.

On Unix-based systems (including Mac OS X and Linux), GIMP is also an “X” application: GTK+ uses the X Window System (sometimes called X11 since most systems use version 11) for drawing to the screen, creating windows, and other low-level operations. So on Mac OS X, Linux, and other Unix systems, you’ll need X in order to run the GIMP. X doesn’t stand for anything except “the next letter after W.” (Seriously! X was the successor to a window system called W which was in use at MIT in 1984 when X was being designed.)

One thing you shouldn’t need is money. The GIMP is free software: free because you can get it without paying, and also because you’re free to redistribute it or modify it.

You *can* buy the GIMP from various sources, and sometimes these versions include extra goodies like an installer or a collection of cool brushes. But most of the time, you’ll do just as well starting with the free version.

### IF THE GIMP IS FREE SOFTWARE, IS IT LEGAL FOR PEOPLE TO SELL IT?

Yes. The GIMP is licensed under the *GNU General Public License* (GPL), which allows anyone to distribute the software, and even to charge money for it.

However, the GPL does insist that further redistribution be allowed (the seller can’t prevent you from making copies and giving them out to your friends). Also, the source code for the program (the human-readable instructions that make the program work) must be included, or at least be made available.

The GPL also requires that you receive a copy of the license itself along with the program, so you can learn about your rights to share the program, modify it, or contribute to the project.

## Where to Find the GIMP

The GIMP project lives at <http://gimp.org>. There, you can find downloadable versions for several operating systems along with instructions on how to get the GIMP for even *more* systems. You can also find out how to build it yourself if you want to see the latest cutting-edge version.

At [gimp.org](http://gimp.org), you’ll also find a substantial collection of documentation. There’s a full user’s manual, a large collection of tutorials, and information on GIMP’s programming interfaces for anyone interested in writing scripts or plug-ins. You can sample a collection of mailing lists and an IRC channel where you can ask questions, volunteer to help, or just trade tips.

## Notes on the OS-specific Installation Appendixes

The Mac and Windows appendixes have a “Quick Start Version” near the beginning. If you’re fairly familiar with your system and software installation, you’ll probably do fine with the Quick Start alone. However, if you hit a snag, just proceed to the more thorough explanation that follows.

Most Linux and BSD users have it pretty easy, and most systems will already have the GIMP installed. Before you try installing, make sure you don’t already have it! If you don’t, Appendix D will get you on the right track.

Time to get started!



# Installing the GIMP on Windows

**C**oaxing the GIMP onto most Windows computers is a very simple two-step process: first download and install GTK+, and then download and install the GIMP. Both are available from the same site.

Everything should go fine if you have Windows NT or later. GIMP 2.2 will work on Windows 98 and ME (though neither platform was officially supported), but the newer versions of GTK+, needed for GIMP 2.4, no longer work on these platforms. Windows 95 users may be able to use GIMP 2.0 (read on for details).

Don't forget to read Appendix F (First Startup) after you're done with the installation. In this appendix, you'll find the following:

- The Quick Start Version
- Getting the GIMP for Windows
- Installing GTK+
- Installing the GIMP
- Cleaning up

## The Quick Start Version

1. First, point your browser to:  
*<http://gimp-win.sourceforge.net/stable.html>*
2. Download "GTK+ 2 for Windows" (this is normally the first download link).
3. While you're there, download "The Gimp for Windows" (the second link).
4. Navigate to your download folder and double-click the "gtk+[version]setup.zip." This will launch an installer. Follow the directions.
5. Double-click "gimp[version]setup.zip" and follow the directions again. It's probably best to let it install a desktop shortcut.
6. Double-click on the GIMP icon to start it (and see Appendix F, "First Startup").

**Note** When you click the download links, you'll be presented with a choice of mirror sites. If this step—or any other—is disorienting, read at least the first parts of “The Longer Version” below.

---

## The Longer Version: Getting the GIMP

The *gimp.org* website has a page on “GIMP for Windows” which discusses various aspects of installation in some detail. More detail than you probably want, really. Referencing this matter can be interesting and informative, but if you're in a hurry, it's probably best to skip right along to the download page.

Like many free and open source projects, GIMP for Windows is hosted on SourceForge, a valuable resource for all computer owners. So point your favorite browser to:

*<http://gimp-win.sourceforge.net/>*

and click on “Stable.”

Normally, this page has two main sections, the “Stable version download” and “Additional packages.” Generally, the latter will not be needed and can safely be ignored.

The first section will have a link to download the most recent versions of GTK+ and the GIMP for Windows. You'll need both.

Click on the link under HTTP (it usually says “Download”) for each program in turn. The order of download is unimportant, but the install order will matter.

When you click the download link, you'll be referred to a page showing *mirror sites* on many continents.

---

**Note** Mirrors are sites that host exact copies of files from a project or application's main site. This allows *gimp.org* and other project sites to “offload” some of the burden of distributing their products. Most major free and open source projects have mirror sites.

---

It's usually best to choose the mirror in the region nearest you. In most browsers, this will initiate the download automatically.

On the mirror sites page, click on the button next to the file size in the “Download” column. There's a natural tendency to click on the mirror's logo in the left-most column (Host), but that will take you to a website rather than delivering the file!

Normally, both downloads add up to about 11MB.

On completion, you should end up with two files in your download folder. One should look like `gtk+[version]setup.zip`, and the other `gimp[version]setup.zip`.

---

**Note** The *[version]* notation is there because the GIMP can be updated at any time. So it's impossible to predict here exactly which version will be on the site by the time you read this. However, this book should be appropriate for all versions of GIMP, starting with 2.

---

Of course, you'll need to unzip the files. You probably already have an unzip utility, but if you don't, there are many free programs available that will work.

Each file should unzip to a folder. Inside each folder will be an installer. On many systems, expanding the zip file will automatically open the folder for you.

## Installing GTK+

Double-click on the installer in the GTK+ folder. It should be named something like `gtk+[version].exe`.

You should see a welcome screen with a colorful GTK logo. After basking in the warm fuzzies, click on the *Next* button.

A copy of the *GNU General Public License* will appear. It's interesting reading, all about your rights regarding free software. You're free to skip it, but you do have to click *Next* again.

Up comes the usual Select Destination dialog. The installer will present you with a default location. Unless you have a very good reason and know exactly what you're doing, you should accept the default and, you guessed it, click *Next* again.

Now you get to select which components get installed. The base component is not optional, but is listed for completeness. GTK-Wimp will make the GIMP look more like your other Windows apps, and is generally a good thing. Translations are not necessary for English speakers (they won't hurt anything besides taking up a little more disk space), but are a good idea for everyone else.

---

**Caution** GTK-Wimp may rarely cause conflicts on some older systems. If you find the GIMP complaining about fonts a lot after installation, you might try reinstalling GTK without GTK-Wimp. So it's a good idea to remember where this screen is.

---

If you leave GTK-Wimp checked, you'll be warned that it requires you to use only TrueType fonts in your Windows display. You probably already are using TrueType—but if you find after installing that the fonts in GIMP windows look ugly, try changing your display fonts (usually *Appearance* in the *Display Properties* window). If that doesn't work, try reinstalling without GTK-Wimp.

Oh, you'll have to dismiss the warning to continue the installation, of course.

Time for the Moment Of Truth: the Ready To Install dialog. This shows a list of the options you've chosen and what will happen when you take the plunge. Go for it! Click *Install*.

After a brief display of The Bar That Moves Right (five seconds on my machine) you get the Finished! dialog, with a button that says *Finish*. Click it and you're done.

If all went well, you now have GTK+ installed on your computer and otherwise nothing is different. It's usually well behaved, and you shouldn't even notice it's there.

Time to get the GIMP working!

## Installing the GIMP

Navigate to the `gimp[version]setup` folder you unzipped earlier. Double-click the installer. It should be named something like `gimp[version].exe`.

Up comes another welcome screen, this time with cute GIMP graphics. After finding Wilber, click on the *Next* button.

Windows 98 and ME users will be treated to a screen telling them the GIMP is no longer actively supported for those platforms. This is mostly true, though GIMP 2.2 versions will probably work. The purpose of this disclaimer is to avoid getting complaints if they don't. Some people who use their spare time to give you free stuff are just so lazy. Click *Next*, of course.

There's that *GNU General Public License* again! It won't have changed since you last saw it, so see how fast you can click the *Next* button.

Déjà vu: the Select Destination dialog again. The installer will present you with a default location. Unless you have a very good reason and know exactly what you're doing, you should accept the default and, you guessed it, click *Next* again.

Time to select which components get installed. The base component is not optional, but it is listed for completeness. Translations are not necessary for English speakers (they won't hurt anything), but are a good idea for everyone else. The Gimp FreeType plug-in is theoretically optional, but if it's there, leave it checked. (It gives you some fun options for adding text to images.)

After clicking another *Next*, you'll be presented with what may be an unfamiliar screen: "Select file associations." If you click on any of the available selections, that kind of file will automatically open in the GIMP from then on. The GIMP is an excellent program, but you're probably not familiar with it yet—so you should almost certainly not select any of these options (thoughtfully, that's the default).

*Next!*

Now you're offered an opportunity to place a GIMP icon in a custom Start Menu folder. You can change the name of the folder, or check the box near the bottom that says *Don't create a Start Menu folder*. The latter is probably the best choice for now, followed by yet another *Next*.

On "Select Additional Tasks" the defaults are perfect. You should probably let the installer make a desktop icon, and a quicklaunch icon is probably overkill until you see how often you want to use the GIMP. Make your choices and click *Next*. Will this ever end?

Yes! You're at the ever-scary "Ready To Install" screen. This shows a list of the options you've chosen and what will happen when you take the plunge. Brace yourself and click *Install*.

After watching the progress bar for about three times as long as GTK+ took, you get the "Completing" screen. Usually, there's a checked option to *Launch The GIMP*. It's probably best to leave it checked and get that over with now, since you already have a good rhythm for clicking buttons. This time it's *Finish*.

---

**Note** If you left the option to launch checked, be sure to refer to Appendix F, "First Startup."

---



## WHAT IF I HAVE WINDOWS 95 (OR EARLIER?)

You're not completely on your own, but you do have to drop back a bit. GIMP 2.2 and later simply won't work right on Windows 95 or earlier. However, you should still be able to run GIMP 2.0, which this book will cover very nicely (the minor differences should cause little trouble). To get it, point your browser to:

*<http://gimp-win.sourceforge.net/old.html>*

There, you'll find older versions of GTK+ and GIMP that worked fine with Windows 95. However, depending on how much you've customized your machine, your mileage may vary. It might be time to think about an upgrade.

## Cleaning Up

Once you've completed the installation, you no longer need the zip files you downloaded, and can safely move them to the recycle bin (and empty it).

You don't need the installer (.exe) files or the folders they're in either. You can delete them, though it's probably best to keep them around for a while in case you need to reinstall for some reason.





# Installing the GIMP on Mac

**D**epending on the state of your system, installing the GIMP can be as simple as downloading a typical Mac application, or a bit more complicated. You may also have to get and install X11.

In this appendix, you'll find the following:

- Requirements
- The Quick Start Version
- Getting the GIMP
- Getting X11
- Installing X11
- Installing the GIMP
- Other GIMPs and X11s for Mac OS X
- Cleaning up

## Requirements

First, GIMP on Macintosh requires OS X. It will not run on earlier versions of Mac OS, largely because the building blocks it needs (GTK+ and X11) were never ported to the earlier Mac systems.


On Mac OS X, the first step to installing GIMP is making sure you have X11. Apple distributes a version of X11, and you may already have it installed: it ships on the OS X installation CDs. If you don't have it and can't find your disks, you can get it from Apple's download site at the same time you get the GIMP.

## The Quick Start Version

If you already know about the Apple OS X software download site, are familiar with typical software installation, have OS X 10.3 or later, and don't need drawing tablet support, here's all you need to know.

1. Go to Apple's download site and get Gimp.app from the "UNIX and Open Source" category.
2. If you haven't already installed X11 and don't have it on CD, get it from the "Apple software" category.
3. Install X11 before you install GIMP. The .dmg includes a self-guiding installer.
4. Open the Gimp-*[version]*.dmg and drag Gimp to your Applications folder.
5. Double-click on the "Turn On X11 Focus Follows Mouse" icon. (For more information, see the note in the "Installing Gimp.app" section.)
6. Double-click on the Gimp icon to start it. It will automatically start X11 if it's not already running, and otherwise act (and mostly look) like any other OS X program. Then see Appendix F, "First Startup."

There is one way Gimp.app will act differently from standard OS X programs, and you may find it confusing. The main menu bar at the top of your screen will show *X11* rather than *The GIMP*, as in Figure C-1. This means all Gimp.app *File* operations will have to be done from menus inside GIMP windows rather than the top menu bar. This is true of all programs running under X11 on Mac OS X—the GIMP is not singled out!



**Figure C-1.** *The strange behavior of the main menu when X11-based programs are running*

## The Longer Version: Getting the GIMP

The *gimp.org* website has a page on "GIMP for MacOSX" which discusses the various Macintosh versions available and where to get them. There are several options, but the easiest for most people is Gimp.app, which includes all the GTK libraries you need as well as the GIMP itself.

---

**Caution** If you are running OS X 10.2 (Jaguar) or earlier, neither Gimp.app nor Apple's X11 will work with your system. You can upgrade your OS X, or skip to "Other GIMPs and X11s for OS X" near the end of this section to find out how you can get the GIMP running on your system.

---

The easiest place to get Gimp.app is from Apple's download site. In the Apple menu in the extreme left of your main menu bar, choose *Mac OS X Software....* Or if you prefer to type in URLs, point your browser to:

<http://www.apple.com/downloads/macosx/>

Near the bottom of the “Categories” column on the left side you’ll find “UNIX & Open Source.” Click on that link.

Somewhere in the list you’ll find Gimp.app—for months it was *the* Featured Download, but by now it has probably been moved. If you click on the link to the description page, you’ll note it requires “Apple’s X11” also. The *Download* button will take you to a page showing various mirror sites in countries all over the world.

Mirrors are sites that host exact copies of files from a project or application’s main site, usually at no charge. This allows *gimp.org* and other project sites to offload some of the burden of distributing their products. Most major free and open source projects have mirror sites.

It’s usually best to choose the mirror nearest you. This will initiate the download automatically in most browsers.

---

**Caution** Click on the button next to the file size in the “Download” column. There’s a natural tendency to click on the mirror’s logo in the leftmost column, but that will take you to a website rather than delivering the file!

---

If you’re on dialup, the download can take quite a while since the file is over 40MB. On completion, you should end up with a file called Gimp-*[version number].dmg* in your download folder.

Stop here. Don’t install it yet, before you’ve checked for X11.

## If You Need to Get X11

How do you know you if you already have X11?

The direct way to check is to look in your Applications folder for “X11.” It should have an “X” icon. If it’s not there, try searching for files named “X11.”

If you don’t find it, you can check your original OS X disks. No luck? Can’t find the disks? It’s probably best to download it now, since you should install it before the GIMP.

Point your browser back to the Apple download pages. It’s fastest to use the *Search Downloads* feature, which is normally in the upper-left, just under the Apple header. Type **x11** into the search bar and make sure the popdown menu has either *Apple* or *All Downloads* selected.

There is no *Go* or *Search* button, so just hit Return on your keyboard. At the top of the resulting list (or near it) will be “X11 for Mac OS X - Apple.” That’s what you’re after.

---

**Note** You may also notice an application called *XFree86* in the “Unix and Open Source” category while you’re looking for X11. See “Other GIMPs and X11s for OS X” near the end of this appendix for further information.

---

Click on the link and follow the same download procedure you used for Gimp.app. However, this time it will come directly from Apple’s servers, so you don’t have to choose a mirror site. Usually X11 and Gimp.app are about the same size.

## WHAT IS THIS X11 THING?

The X Window System has been around for ages. Apple calls it “X11” but you may also have heard of X11R6, Xfree86, X-windows, the X Window System, or even just X. The last is particularly confusing for Mac users since Apple’s current OS is called, well . . . it’s called “ten,” but it *looks* like X!

X is basically the graphical part of most Linux and Unix distributions. It plays the same role on these machines that Quartz and Aqua play on OS X: it handles graphics, windows, and the input from the keyboard and mouse. Meanwhile, the heart of the operating system, called the “kernel” (in the OS X case, a combination of the Mach and FreeBSD Unix systems, called Darwin) handles non-graphical tasks such as the network and the file system.

Apple offers X for Mac users who would like to take advantage of the enormous collection of Linux, Unix, and GNU programs designed for it. The GIMP is one of those programs.

## If You Need to Install X11

If you have system disks with X11 on them, double-click on the icon and you’ll be presented with an installer. If you downloaded it, you’ll find X11User.dmg in your download folder, which you can double-click; it will expand into a disk image. Inside, you’ll find an icon called X11User.pkg. Double-click that and you’ll be presented with a standard OS X installer.

From there on it’s the usual sequence: agree to the license, choose a place to install, click a few times, and you’re done.

No further action should be required regarding X11 for you to run Gimp.app. If X11 is not already running when you start GIMP, the launcher will take care of that for you.

---

**Note** You may have noticed references to Gimp.app, GIMP, and the GIMP. Why the different usages? The authors of the GIMP tend to capitalize it because it’s an acronym for the GNU Image Manipulation Program. Sometimes they call it *the* GIMP, and sometimes just GIMP. When it was ported to Mac OS X, the capitalization was not used. Also in keeping with typical OS X style, the packaged program was called Gimp.app—which means Mac people are likely to see Gimp, GIMP, The GIMP, and Gimp.app used interchangeably.

---

## Installing Gimp.app

Double-clicking on the Gimp-*[version]*.dmg in your download folder will expand a disk image called, appropriately, Gimp. Inside you’ll find a collection of files including one named Gimp.

At this point, you can simply drag the icon labeled Gimp into your Applications folder (or wherever you like to keep your downloaded programs).

You’ll also note the Get X11 for OS X 10.*[version]*. This did not work on my copy; I don’t suggest counting on it.

LICENSE.txt will have a copy of the *GNU General Public License*, which you are not required to read just to use the program. It’s interesting, but you’re free to skip it. It won’t be on the test.

The README.txt file will have very up-to-date information about the version you just downloaded. You’ll find longer and more specific installation instructions tailored to that version. There are also some comments about the people who made this program for you.



Especially interesting is an explanation of the final item in the folder: “Turn On X11 Focus Follows Mouse.” Use this option! If you don’t, you’ll have to click twice on each button you use in the GIMP. Double-clicking on that icon in the Gimp installation folder cures the problem.

If you decide later that you don’t like “Focus Follows Mouse,” open Terminal.app and type:

```
defaults write com.apple.x11 wm_ffm false
```

Hit Return at the end of the line, of course.

## Other GIMPs and X11s for OS X

For a typical OS X 10.3 (Panther) or 10.4 (Tiger) user, this information is probably just “background.” Your best bet is to install Gimp.app and Apple’s X11.

But if you are running OS X 10.2 (Jaguar) or 10.1 (Puma), you *can’t* run either Gimp.app or Apple’s X11, and must rely on other resources. Also, anyone who wants to use a graphics tablet or other specialized input device may need a different X11 implementation. Apple’s X11 does not generally support such hardware. Try it with Apple’s X11 first; but if it doesn’t work, don’t give up until you try a different X.

(If you’re running OS X 10.0—Cheetah—I know of no way to run the GIMP short of porting it yourself.)

First you need a compatible version of X11. There are various versions available; some cost money, others are free. Probably your best choice is XFree86, which is free and compatible with OS X 10.1 or later. It’s a completely open source implementation that is generally more up-to-date than Apple’s version. The downside is that it usually runs a bit slower.

You can download it from various places, including Apple’s own site, under the name XFree86. Or you can shop direct at <http://www.osxgnu.org>.

You can also get XFree86 via the *fink* or *darwinports* methods (see below).

There’s also XDarwin, a pay application (not free), which works with some graphics tablets.

Meanwhile, OS X 10.1 and 10.2 users will also need a compatible version of the GIMP.

A good option is to install Fink (mentioned in the Gimp.app README.txt). It’s a very extensive collection of software along with the tools to install and use it. You can look it over at <http://fink.sourceforge.net>. Once you have Fink, you can install XFree86 *and* the GIMP by simply typing into Terminal.app:

```
apt-get install xfree86 gimp
```

Alternately, for OS X 10.2 or later, you might wish to compile (build) it on your own system. This would normally be a bit daunting, but DarwinPorts comes to the rescue!

You can find out more about DarwinPorts at the Apple download site, again under the Unix and Open Source category. Using its automated tools, you can build and install your own GIMP. All you have to do is type:

```
port install gimp
```

DarwinPorts will do the rest, which may also include building various dependencies. This can take quite a while (several hours).

All these methods work fine, and both Fink and DarwinPorts also supply different versions of the GIMP.

If you'd like to “shop direct” for the GIMP by going to the original developers' site, point your browser to <http://gimp.org/macintosh/>. What you'll find there is . . . Gimp.app! But there's also a set of links that can eventually solve just about any weird problem you might have.

Plus, there is some interesting background information along with the latest notes. Downloading Gimp.app is essentially no different from using Apple's site—you will be presented with the same list of mirrors eventually.

## Cleaning Up

Once you've installed X11, you no longer need X11User.dmg or Gimp-*[version]*.dmg. You can delete them, though it's probably best to keep them around for a while in case you need to reinstall for some reason.

You launch the GIMP just like any other OS X program—go to your Applications folder (or wherever you chose to drag and drop it) and double-click the Gimp icon.

The first time you start the GIMP, an icon featuring Wilber (the mascot) will appear in the Dock in the usual place. Like other OS X applications, if you move the icon to another location in the Dock, it will stick. From then on, you can launch the GIMP just by clicking its Dock icon. If X11 is not yet running, the Dock launcher will take care of that for you, just as if you had launched from the Applications folder.

You're ready to start exploring the GIMP! Don't forget to read Appendix F for some tips on “First Startup.”



# Installing the GIMP on Linux or Unix

**O**n Linux, installing the GIMP is usually very easy. Nearly all Linux distributions already include GIMP, either on the installation CDs or as an extra you can download using built-in software installation tools.

In this appendix, you'll find installation notes for the following:

- Red Hat/Fedora Linux
- SuSE Linux
- Debian (including variants such as Ubuntu or Libranet)
- Linspire
- Mandriva
- FreeBSD and NetBSD
- Other Unix variants

---

**Caution** If your Linux distribution isn't current, it may not yet have GIMP 2.4. Most of this book should make sense for GIMP 2.2 users and even GIMP 2.0 users; but there are a few features that weren't added until 2.4. If you're eager to try out the latest and greatest, skip ahead to Appendix E, "Building from Source."

---

After installing, don't forget to glance through Appendix E, "First Startup."

## Installation Notes for Various Distributions

Each Linux distribution has a slightly different set of tools for installing software. With any luck, you're already familiar with how to install packages on your system, but here are some guidelines for the most common Linux distributions.

## Red Hat

If GIMP wasn't already included when you installed your system, you can probably get it using the software installation window. If you have installation disks, the GIMP is probably on them. Alternately, from a terminal (as root), you can type the following:

```
up2date gimp
```

Fedora users, use this command:

```
yum install gimp
```

## SuSE

SuSE normally installs the GIMP by default. The GIMP is always present on all versions of the install disks, but it may not be the most current build available.

With that in mind, it may be a good idea to use SuSE's usual YaST software screen to install the newest version.

## Debian Variants

Any Debian-based system should be able to install the GIMP by typing (as root):

```
aptitude install gimp
```

This will apply to Ubuntu, Libranet, Knoppix, and a host of others. You can see a mostly complete list of Debian flavors at:

<http://www.debian.org/misc/children-distros>

Alternately, you can use a graphical program such as *synaptic* (or whatever package manager is available in your desktop menus) to choose and install the GIMP from a menu of software.

## Linspire

The GIMP is part of the Click-N-Run Warehouse, and that may be the easiest way to get it. If you don't have Click-N-Run membership, you may be able to get the GIMP using the Debian tools *aptitude* or *apt-get* (see "Debian Variants," above).

## Mandriva

The GIMP is available through RPMDrake, Mandriva's software installation tool. If you need a version newer than the one they offer, you may be able to find an RPM from Mandriva's "cooker" or other Mandriva sites.

## FreeBSD and NetBSD

On FreeBSD, it's easy. The GIMP has been faithfully ported to all FreeBSD Unix versions for years, though it is typically a little behind the current release.

FreeBSD has a number of ways to acquire and install software. The GIMP may be on the set of installation disks you used. If not, you can use the ports system if you're familiar with it, or type the command:

```
pkg_add -r gimp
```

On NetBSD, it's slightly trickier: you need to find the NetBSD GIMP package online, and then either download it and pass it to *pkg\_add*, or run **pkg\_add url** where *url* is the web address of the package.

The BSD variant OpenBSD has a GIMP package available, but it's so old that you're probably better off building from source (see Appendix E).

## Other Unix Versions

Many proprietary Unix versions from companies such as Sun and SGI have freeware or open source sites that provide precompiled versions of popular software such as the GIMP.

Unfortunately, such sites tend to be way out of date. SGI's freeware page offers GIMP 1.2.5. Sun's is even further out of date, offering 1.2.1—released April 2002! Fortunately, there are sometimes other sources for downloads of newer versions, especially for Solaris. Check the main *gimp.org* site for an up-to-date list.

If you can't find a build for your platform, though, move on to Appendix E, "Building from Source."







# Building from Source

**F**or really uncommon platforms, the only answer may be to build the GIMP yourself. And there are other good reasons to build your own GIMP.

Perhaps you want the absolute latest and greatest version, to check out the hot new features. Maybe you want to have several versions installed at once, in different places. What if you're running an older OS, and the only GIMP available for it is too old? You might want to contribute to the GIMP project, or just read the code and play with it a bit to understand how it works. These are all good reasons.

In this appendix, you'll find the following:

- Getting the source
- Notes on assembling dependencies
- The three stages of building: configure, make, make install

## Getting the Source

GIMP source usually comes packaged as a “tarball,” a compressed file named `gimp-x.y.z.tar.gz` or `gimp-x.y.z.tar.bz2`. Unpack this with the tar program: **`tar xzvf gimp-x.y.z.tar.gz`** or **`tar xjvf gimp-x.y.z.tar.bz2`**.

If you want a very recent version of the GIMP, you may be better off checking out a copy from CVS, the Concurrent Version System. Follow the CVS instructions at [gimp.org](http://gimp.org) to check out a CVS copy of the source.

Once the source has been unpacked, it's a good idea to peruse the files `README` and `INSTALL`. The latter is especially important, as it contains advice on how to build the GIMP.

## ASSEMBLING DEPENDENCIES

Building the GIMP is easy. However, assembling the *requirements* to build it can be difficult. GTK+, the underpinnings of the GIMP, needs a lot of additional libraries to work properly. The library list includes names such as GLib, Pango, ATK, FreeType2, Fontconfig, and libart2.

The *gimp.org* page “GIMP from Source” has a requirements list, but it’s slightly misleading—the link for each library points to that library’s source code. You may not need to build each of these libraries. Try first to find pre-built packages for your system. Building them is a last resort.

If you do find pre-built versions of the required libraries, make sure you install the “development” components of the libraries as well as the runtime.

On Linux, the development components are usually packaged separately. Typically, they’re named something like *[package name]-dev* or *[package name]-devel*. For instance, on Debian, in addition to installing *libgtk2.0-0* you should also install *libgtk2.0-dev*.

You will also need a C compiler, since the GIMP is written in the C programming language. On Linux or Unix, this is usually just a matter of installing the *gcc* package to get the GNU C compiler. On Mac, *gcc* is on the Development Tools CD which came with your copy of OS X. Windows users can download *gcc* from [www.cygwin.com](http://www.cygwin.com), but if you don’t expect to be doing heavy software development, you may be better off with a smaller package containing the minimum you need to compile software: the Minimalist GNU for Windows, available at <http://www.mingw.org> (you will need both MinGW and MSYS).

Finally, during the build process you may find that you need other packages too (see below, “Configure”). Try not to get discouraged; it can take a while to assemble all the dependencies, but once you get them all together, you’re done and you’ll be able to build the GIMP and lots of other interesting software as well.

Once you have a compiler and all the required libraries, it’s time to start the build.

You will need a terminal open in order to build GIMP. All commands in this section are typed at your terminal prompt. It’s not important what terminal window you use; on Mac you can use either Apple’s Terminal program or the *xterm* program that comes with X11; on Windows, either the DOS shell or one of the terminal programs available from Cygwin; on Linux, any terminal program will work.

Most open source software, including the GIMP, is built by typing three commands, one after the other:

```
./configure  
make  
make install
```

### Configure

The *configure* step of the build is usually the hardest. This is where GIMP’s build system searches your computer for the dependencies it needs. If it fails to find something, it will stop with an error message. Interpreting these error messages can sometimes be challenging.

---

**Caution** If you're trying to build a cutting-edge GIMP from CVS, type `./autogen.sh` instead of `./configure`. You'll probably also need recent versions of the packages `autoconf` and `automake`.

---

Sometimes the messages are straightforward: it wants version 2.8.4 of something; you only have version 2.6, so the solution is to upgrade.

Other times it's not so easy. If `configure` says it needs a library and you're fairly sure you have it, check where the library is installed. Perhaps it's not in a standard location, or you need to update the system to notice the newly installed library. Or maybe you have the “runtime” part of the library installed, but you don't have the “development” part (usually this consists mostly of “C header files,” small text files ending in an `.h` extension).

You can customize some aspects of the GIMP by running `configure` with arguments: `./configure option1 option2 ...`. Type `./configure -h` (that's h for “help”) to get a list of options. The default—no extra options—should produce a usable GIMP on most platforms.

If you need to install more than one version of GIMP on the same system, use a different “--prefix” (that's two minuses) so that each GIMP will look in a different place for its files. For instance:

```
./configure --prefix=/usr/local/gimp-2.4.2
```

The `configure` program will also build extensions to the GIMP if you have the appropriate libraries. For instance, support for many different image formats requires additional libraries. If you don't have `libjpeg` installed, you may end up with a GIMP that works fine but can't read or write JPEG files. Watch the messages printed by `configure` as it does its work: it will tell you if there are components it doesn't find. If they are components you care about, you can look for the appropriate libraries, install them, and then run `configure` again.

If you hit a wall trying to interpret `configure`'s messages, try looking at the file named `config.log`. It contains all of `configure`'s output, starting with the arguments you passed to `configure`. The useful information, about what `configure` was looking for and how it failed, will usually be at the very end of the file. `Configure` does some fairly esoteric things so it may be difficult to interpret the file, but sometimes it's fairly clear what's wrong. If all else fails, you can provide the errors at the end of `config.log` to a guru on the mailing lists or IRC channel to get help. (Please don't paste long error messages into an IRC channel without asking first, though!)

## Make

Once you've configured, the *make* stage is time-consuming, but easy. Simply type **make**. If all goes well, you will see a huge amount of output scrolling down your terminal window. Unless you have a very fast machine, this will probably take quite a while (half an hour or more). Find something else to do while GIMP is compiling, or go have a snack.

When the output stops scrolling and you get your prompt back, it's done. If there's no error message, you've successfully built GIMP, and you're ready to install it.

## Make Install

The *make install* phase is how you install the GIMP after you've built it. If you're installing to a system directory, as is usually the case (though you can override that with `configure --prefix`, as discussed above), you'll probably need "root" or administrator privileges on your machine for this stage.

Then type **make install** and watch the messages scroll down. This will produce nearly as much output as the `make` phase did, but it will take a lot less time.

When `make install` finishes, you're ready to try out your new GIMP!

---

**Note** On some Linux systems, you may see error messages about libraries not found when you first run GIMP compiled from source. If that happens, try running the command `ldconfig` as root.

---



# First Startup

**T**he first time you launch the GIMP, you'll see a screen with a description of free software and a disclaimer of warranty, plus the very important *Continue* button. Click it (after reading the text, if you haven't already).

---

**Note** Let's ruin the plot right now. Your best move is to click the *Continue* button every time you see it until they stop letting you. Works every time!

---

Next up is an incredibly complicated Personal GIMP Folder screen. Don't even think about reading it. Just click *Continue*.

The GIMP will spew lines and lines listing the files installed. Don't worry, they're harmless. (In the unlikely event this stage fails, check permissions on the folders it was trying to use.) And just to keep you in the loop, you get to click *Continue* again.

Now you're offered the opportunity to fine-tune the Tile Cache and Swap File, about which you probably know nothing. I guarantee you don't want to know the algorithms involved. Your only rational choice here is to click *Continue*.

---

**Note** All the settings you just agreed to are available in the Preferences window, if you decide later that you want to go back and change them. Chapter 1 has a discussion of GIMP's preferences, and Chapter 12 has more details on what's going on behind the scenes. You may eventually want to increase the Tile Cache size if you have a lot of RAM and work with large images.

---

After you do, the GIMP will actually start! Following an impressive progress-bar display (don't worry, it doesn't *always* take as long to start up as it does the first time), it will grace you with the GIMP Tip of the Day. The tips are helpful, but if you get tired of them, pay close attention to the checkbox in the lower-left. (You can get them back via GIMP's Help menu.)

You're done. The GIMP is yours! Bend it to your will.

Well, there's one last thing: clicking the Close button in the Main window—the one titled “The GIMP” with all the tool icons, commonly called the Toolbox window—will dismiss all other windows and exit the application. Just so you know.





# Index

## ■ Symbols

- # (hash) comment character, 439
- #! characters (shebang), 439
- #include, 446
- \$ (dollar sign), 444
- \* (wildcard character), 419
- \*/ sequence, 446
- [ (single open bracket), 444
- [] (square brackets), 441
- + button, 322

## ■ A

- Acquire menu, 465–466
- Activate the focused image option, 22
- active layer, 83
- Adaptive filter, 259
- Adaptive Hybrid, 464
- Adaptive supersampling, 146, 263
- Add button, 28
- Add CC field, 480
- Add Fog filter, 285
- Add Layer Mask dialog, 206
- Add mode, 191
- Addition mode, 343–345, 399–400, 406
- additive colors, 298–299
- additive stacking, 406
- Adjust Output. button, 463
- Adjustment option, Ink Pen, 132
- ADJUSTMENT widget, 437–438
- Advanced area, Save as JPEG dialog, 38
- Advanced Options, GIMPpressionist, 285
- Advanced Options tab, New Image dialog, 118
- Affect option, Rotate tool, 61
- Affect tabs, 97
- Affected Area option, 140
- Airbrush tool, 130
- Alignment option, 236
- All feature, Select menu, 169
- All Layers option, 163
- Alpha channel, 52, 134, 251, 428
- Alpha to Logo scripts, 292–293
- Ambient option, Bump Map, 280
- Amount parameter, Unsharp Mask, 64
- Amount slider, 270
- Anchor button, 99, 146
- animated brushes, 370–374

- animation, 34, 41, 292
- Anti erase option, 135
- antialiasing, 81, 129, 137, 173, 216, 260, 353
- aperture, lens, 244
- Appearance, Display Properties window, 489
- Apple OS X software download site, 493
- Applications folder, 496
- Apply button, GIMPpressionist, 276
- Apply Canvas filter, 273, 276
- Apply Lens effect, 268
- Apply preset, GIMPpressionist, 277
- apt-get tool, Debian, 500
- aptitude tool, Debian, 500
- arithmetic operations, 431
- Artistic filters, 273–279
- Artistic menu, 275
- aspect ratio, 29
- astrophotography, 404
- astrovideographers, 409
- Auto button, Levels dialog, 51
- Auto menu, 57
- Auto option, 462
- Auto Shrink button, 44
- Autocrop, 445
- Autocrop All Layers function, 454
- averaging, increasing resolution by, 409
- Azimuth filter, 260
- Azimuth option, Bump Map, 280

## ■ B

- background color (BG), 85, 140, 267, 427
- background layer, 163, 385, 390
- Behind drawing tool mode, 344
- Behind mode, 128
- bevels, 267
- Bezier Paths
  - adding nodes or segments and moving paths, 185
  - curved paths, 184–185
  - defining a path, 175–177
  - moving or modifying an existing path, 185–186
  - overview, 174–175
  - Paths dialog, 177–178, 180, 182
- bg-color parameter, 436
- Bilinear option, Interactive Warp filter, 263

- bits per channel, 300
- bits per pixel (bpp), 300
- black and white. *See* grayscale or black and white
- Black level control, 259
- Black Pullout parameter, 305
- black velvet paintings, 393–394
- Blast filter, 266
- blending, 292, 341, 435
- Blinds filter, 259
- blue sphere, 455
- blurring, 90, 242, 246, 257–259, 338
- BMP (Windows Bitmap format), 36
- bookmarks, 28
- Border feature, Select menu, 170
- Border filter, 264–265
- borders, fuzzy, 334–340
- boundaries, when cropping, 456–457
- boundary selection, 171
- bpp (bits per pixel), 300
- brightening, 222
- brightening and darkening photos
  - brightness-contrast, 44–47
  - curves, 52–54
  - levels, 47–52
  - other exposure adjustments, 54, 56–57
  - overview, 44
- brightness, methods of measuring, 307
- Brightness slider, 70
- Browse for other folders option, 33
- brushes
  - making
    - animated brushes, 370–374
    - importing animated brushes, 374
    - overview, 368
    - simple image brushes, 368–370
    - overview, 122–126
    - parametric vs. image brushes, 126
- BSD variant, 501
- Bucket Fill tool, 139–142, 469
- buffer allocation, 456
- bug-tracking program, 479
- Bugzilla page, 479
- built-in patterns, 376
- Bump Map filter, 260, 279–280
- Burn mode, 345–346
  - to cut through haze, 387–390
  - making eerie colors with, 402–403
- Burn-in option, 292
- Burning Well, 482
- Buttons, Toolbox's Xtns menu, 294
- By Color feature, Select menu, 169

## C

- C compiler, 504
- C function, 446
- C plug-ins, 448
  - building, 421–422, 464
  - installing, 420–421, 463
  - writing
    - handling edges, 457–458
    - layer coordinates vs. image coordinates, 455–456
    - making new plug-in, 454–455
    - overview, 444–445
    - pixel regions, 450–454
    - routine that does the work, 449–450
    - run procedure, 447–449
    - searching image for blank rows and columns, 452–453
    - setting boundaries, 456–457
    - tiles and the tile cache, 451–452
    - zealous crop, 445–447
- cadr (Contents of Decrement Register), 432
- Calibrate chain link icon, 21
- camera lenses, 410
- Camouflage, Toolbox's Xtns menu, 294
- Cancel button, 6
- canvas size, increasing, 161–163
- Cap style, 137
- Cartoon filter, 274
- Center buttons, 463
- chain link, 19, 30, 93
- ChangeLog, 476
- Channel Attributes, 199
- channel masks, 197–199, 201
- Channel menu, 251
- Channel Mixer, 311
- channel preview size option, 471
- Channel to Selection, 201
- channels, 52
- Checkerboard pattern, 286
- Clear option, Text Editor window, 79
- Click-N-Run Warehouse, 500
- Clip result option, Rotate tool, 61
- Close button, Main window, 507
- Clothify filter, 273
- Clouds, 285–286
- CML Explorer pattern, 286
- CMY colors. *See* RGB and CMY colors
- CMYK colors, 303–304
- CMYK tab, 305
- color channels, 323–324
- color correction, 52
- color depth, 300–301
- Color Enhance function, Auto menu, 57
- Color erase, 128
- Color Erase drawing tool mode, 344
- Color index, 322
- color management, 21
- Color Management category, 330
- Color mode, 347

- Color option, 90
- Color Picker tool, 322–323
- Color preset, GIMPpressionist, 277
- Color Sensitivity expander, 216
- colored pencil drawings, 391–393
- Colorify filter, 382
- Colorify plug-in, 313–314
- Colorize, 314
- colorizing images, 381–384
- colormap editor, 321
- colorrrc file, 473
- colors
  - alternative color choosers, 304–306
    - CMYK tab, 305
    - triangle color selector, 305
    - Watercolor selector, 306
  - CMYK colors, 303–304
  - color channels, 323–324
  - color decomposition, 324–329
    - decomposing to HSV, 327–329
    - overview, 324–325
    - QuickMask, 326–327
    - Threshold tool, 326–327
  - Color Picker tool, 322–323
  - color profiles, 330
  - coloring monochrome images and making sepia photos, 311–315
    - automatic conversion with Old Photo filter, 312–313
    - manual conversion for fine control, 313–315
  - Hue, Saturation, and Value (HSV) color model, 302–303
  - indexed color, 317–322
    - converting to indexed mode, 318–319
    - dithering an indexed image, 319–320
    - overview, 317–318
    - redesigning for better indexed results, 320–322
  - profiles, 330
  - RGB and CMY colors, 298–302
    - additive colors, 298–299
    - color depth, 300–301
    - indexed palette, 301
    - representing RGB color in web pages, 301–302
    - subtractive colors, 299
  - Threshold tool, 315–317
  - working with grayscale or black and white, 306–311
    - Channel Mixer, 311
    - Decompose dialog, 309–310
    - Desaturate function, 308
    - Grayscale mode, 308
    - Hue-Saturation dialog, Colors menu, 308
    - methods of measuring brightness, 307
- Colors filter operations, 256
- Colors menu, 9, 44, 54, 256, 308, 418
- Colors tab, Fractal Explorer, 288
- Colorspace selector, 19
- colours\_equal routine, 453
- Combine filters, 272–273
- Command Prompt program, 421
- comments, in Lisp, 427
- Compensate for darkening option, Bump Map, 280
- compositing
  - colorizing images, 381–382, 384
  - defined, 381
  - overview, 381
  - self-compositing with layer modes
    - overview, 384–385
    - using Overlay or Burn to cut through haze, 387–390
    - using Overlay or Hard light when light is flat, 386–387
    - using Screen mode for dark images, 385–386
  - self-compositing with modifications
    - adding blurs and other tricks for artistic effects, 395
    - Layer Offset, 390–394
    - overview, 390
  - stacking images
    - increasing contrast by multiplicative stacking, 407–408
    - increasing light by additive stacking, 406
    - increasing resolution by averaging, 409
    - loading all images as layers, 405
    - overview, 404–405
    - reducing noise, 405
    - registering images using Difference mode, 405–406
  - stitching panoramas
    - add rest of images, 415
    - adjust layer mask, 414
    - adjust with other tools if necessary, 415
    - calculate your expected new image size, 411
    - decide on resolution, 411
    - final adjustments, 415
    - fine-tune position of second image, 413
    - load first two images, 412
    - make Gradient layer mask, 412
    - overview, 409–410
    - shooting images, 410–411
  - unrelated images
    - making eerie colors with Burn, 402–403
    - overview, 396
    - using Addition to complement light and dark, 399–400

- using Grain merge to add texture, 403
    - using Overlay for dark images, 397–398
    - using Screen to get lighter effect, 399
    - using Soft light for combining images, 396
    - using Subtract to make cut-out mask, 400–402
  - compositing unrelated images
    - making eerie colors with Burn, 402–403
    - overview, 396
    - using Addition to complement light and dark, 399–400
    - using Grain merge to add texture, 403
    - using Overlay for dark images, 397–398
    - using Screen to get lighter effect, 399
    - using Soft light for combining images, 396
    - using Subtract to make cut-out mask, 400–402
  - compression, 34–35, 41
  - Concurrent Version System (CVS), 476, 481, 503
  - conditional expressions, 431
  - configuration files
    - folders inside GIMP profile, 475
    - gimprc
      - default brush, pattern, palette, gradient, and font, 469
      - default threshold, 469
      - Help browser, 470
      - help locales, 470
      - overview, 468
      - perfect mouse setting, 469
      - plug-in paths, 470–471
      - restoring keyboard shortcuts, 470
      - saving modified images, 470
      - share palette, 469
      - stingy memory use, 469
      - tear-off menus, 470
    - gtkrc, 471–472
    - menu entries and key bindings, 473
    - menurc, 473
    - overview, 467
    - sessionrc, 472
    - window positions, 472
  - Configure Extended Input Devices option, 22, 133
  - Configure Keyboard Shortcuts dialog, 473
  - configure phase of build, 504–505
  - configure program, 505
  - Constraints option, Rotate tool, 61
  - Contents of Decrement Register (cadr), 432
  - context menus, 7–8
  - Continue button, 507
  - contrast, increasing by multiplicative stacking, 407–408
  - control points, 53, 174
  - Control-click, 177
  - controllerrc file, 473
  - Convert Image, 40
  - convolution, 242
  - Convolution Matrix filter, 272
  - Convolve tool, 240
  - Copy filter, 259
  - Coupled Map Lattices pattern, 286
  - Create path from text buttons option, 86
  - Create path from text option, 81
  - Create Selection from Path button, 178
  - Creative Commons, 482
  - crop function, 418
  - Crop tool button, 42
  - cropping, 42
  - Cubism filter, 274
  - Cubism preset, GIMPpressionist, 276
  - Cumulative layers option, 114
  - Current Layer Only option, 44
  - Curve Bend filter, 259
  - curves. *See* lines and curves
  - Curves dialog, 52–54
  - customizable settings
    - color management, 21
    - default grid, 20
    - default image, 18–19
    - display, 21
    - environment, 13–14
    - folders, 23
    - help system, 16
    - image windows, 20–21
    - input devices, 22
    - overview, 13
    - theme, 15–16
    - tool options, 16–18
    - toolbox, 18
    - user interface and keyboard shortcuts, 14–15
    - window management, 22
  - cut-out mask, 400–402
  - CVS (Concurrent Version System), 476, 481, 503
- D**
- dark images
    - using for Screen mode, 385–386
    - using Overlay for, 397–398
  - darkening photos. *See* brightening and darkening photos
  - DarwinPorts, 497
  - darwinports method, 497
  - Dash pattern, 137
  - DCT method control, 38
  - Debian tools, 500

- Debian variants, installing GIMP on, 500
- Decompose dialog, 309–310
- Decompose function, 304
- decomposition, color, 324–329
  - decomposing to HSV, 327–329
  - overview, 324–325
  - QuickMask, 326–327
  - Threshold tool, 326–327
- decompression, 35
- Decor menu, 293–294
- Deep Field project, 404
- default gradient, 144
- default grid, 20
- default image, 18–19, 118
- Default layer option, 471
- default threshold, `gimprc`, 469
- default-font line, 468
- Deform amount option, Interactive Warp filter, 263
- Deform Modes, Interactive Warp filter, 262
- Deform radius option, Interactive Warp filter, 263
- Deinterlace filter, 259
- Delay between frames option, 114
- Delete brush, 126
- Density option, 464
- Depth controls, 260
- depth maps, 272
- Depth Merge filter, 272–273
- Depth option, Bump Map, 280
- Desaturate function, 308
- desaturate images, 56
- Desaturating control, 69
- Description option, 373
- Despeckle filter, 259
- Destripe filter, 259
- Developer FAQ, 476
- development components, 504
- Development Tools CD, 504
- development version, 475
- Device Status dialog, 133
- devicerc file, 473
- dheight, 457
- dialog pops up, 428
- dialogs, grouping by docking, 10–12
- Dialogs menu, 9
- Difference mode, 145, 344–345, 393–394, 405–406
- Difference of Gaussians, 270
- Diffraction Patterns pattern, 286
- digital photographs, and exposure latitude, 220
- digits parameter, 438
- Dilate filter, 272
- Dimension option, 373
- Direction filter, 266
- Directional lighting type, 267
- DIRNAME, 438
- dirty parameter, 450
- Discard Text Information button, 106
- Displace filter, 280–283
- Display preferences, 21
- Dissolve mode, 128
- Distorts filters, 259–266
- Dither Algorithm, 464
- Dither size, GIMPressionist, 285
- dithering, 40, 319–320
- Dithering mode, drawing tool, 146
- Divide mode, 345, 391–393
- Divisions pattern, 286
- `do_zcrop` function, 449
- Dock launcher, GIMP, 498
- docking, grouping dialogs by, 10–12
- Docking menu, 102
- Document History option, 28
- documents file, 474
- Dodge mode, 345–346
- dollar sign (\$), 444
- Dot for dot by default controls, 20
- Download button, 495
- dragging, 43
- DRAWABLE, 439, 481
- drawable dimensions, 457
- drawable layer, 435
- drawable parameter, 450
- drawing
  - adding reflections and shading, 363–368
  - drawing project
    - final touch-ups, 164–166
    - increasing canvas size, 161–163
    - making box using perspective transformation, 153–161
    - overview, 147–153
    - preparing planter box, 161
- Eraser tool, 134–135
- layer modes
  - Addition mode, 344–345
  - Burn mode, 345–346
  - Color mode, 347
  - combining layer modes, 351–356
  - creating depth with, 348–356
  - Difference mode, 344–345
  - Divide mode, 345
  - Dodge mode, 345–346
  - Grain extract mode, 347
  - Grain merge mode, 347
  - Hard light mode, 346
  - Hue mode, 347
  - Lighten only mode, 347
  - Multiply mode, 345

- Overlay mode, 345–346
  - overview, 341–344
  - Saturation mode, 347
  - Screen mode, 345–346
  - Soft light mode, 346
  - Subtract mode, 344–345
  - Value mode, 347
  - lines and curves
    - Airbrush tool, 130
    - brushes, 122–126
    - drawing fuzzy or smooth lines, 128–129
    - drawing tool options, 127–128
    - hard edged lines, 121–122
    - Ink Pen Tool, 131–133
    - overview, 121
    - Pencil tool, 121–122
  - making brushes
    - animated brushes, 370–374
    - importing animated brushes, 374
    - overview, 368
    - simple image brushes, 368–370
  - making gradients, 377–378
  - making patterns, 374–376
  - mask tricks
    - fuzzy borders, 334–340
    - interpreting mask border color, 340
    - overview, 333
    - text fade out, 333–334
  - new image, 117–119
  - overview, 117, 333
  - realism and multipoint perspective
    - overview, 360
    - single-point perspective, 361
    - two-point perspective, 361–363
  - realistic shadows
    - overview, 357–358
    - transparency, 359–360
    - using object to cast own shadow, 358–359
  - shapes
    - Bucket Fill tool, 139–142
    - defining regions, 135–136
    - filling regions, 138–139
    - free select with Lasso tool, 138
    - Gradients tool, 143–146
    - outlining selections, 136–137
    - overview, 135
    - patterns, 142–143
    - stroking, 136–137
    - using layers, 119–120
  - drawing mode, 435
  - Drop Shadow and Bevel, 267
  - drop shadow layer, 94
  - Drop shadow option, GIMPpressionist, 277
  - Drop-Shadow dialog, 102
  - Duplicate brush, 126
  - Duplicate Layer option, 109
  - Duplicate option, 178
- E**
- edge-detection filters, 270–271
  - Edit brush, 125–126
  - Edit Channel Attributes, 201
  - Edit Layer Mask, 207
  - Edit menu, image window, 8
  - Edit mode toggle button, Bezier path, 184
  - eerie colors, 402–403
  - effects. *See* filters
  - Elevation filter, 260
  - Elevation option, Bump Map, 280
  - Ellipse Select tool, 96, 135, 285
  - Ellipse Selection tool, 70
  - embossing, 260, 394
  - Engrave filter, 261
  - Enhance filters, 259
  - ENUM widget, 438
  - environment, 13–14
  - Environment Map option, 267
  - Equalize function, Auto menu, 57
  - Erase Every Other Row filter, 261
  - Eraser tool, 134–135
  - erasing transparency, 135
  - Erode filter, 272
  - Exchangeable Image File Format (EXIF), 38
  - .exe (installer) files, 491
  - EXIF (Exchangeable Image File Format), 38
  - exposure. *See* brightening and darkening photos
  - exposure latitude, 220
  - Exposure setting, 224
  - eye icon, 257
  - eyedropper buttons, 48
- F**
- Fade out, 127, 222
  - FAQ (Frequently Asked Questions) documents, 477
  - Feather dialog, 182
  - feathering, 137, 169, 173, 180, 216, 334, 337, 364
  - Featured Download, 495
  - Fedora, 500
  - FG to BG (RGB), 145
  - File operations, 3, 494
  - FILENAME, 438, 480
  - files
    - formats
      - GIF, 34
      - JPEG, 34



- other formats, 36
    - overview, 33
    - PNG, 35
    - TIFF, 35
    - XCF, 35
  - hidden, 369
  - opening, 27–28
  - saving, 32–33
  - filesystem path, 28
  - Fill Opacity button, 201
  - Fill transparent areas option, 142
  - Fill with BG Color option, 138, 163
  - Fill with Pattern option, 138
  - filling regions, 138–139
  - Filmstrip filter, 272
  - filters, 9
    - adding patterns to layer
      - Clouds, 285–286
      - Fractal Explorer, 288
      - Gfig, 289
      - Nature, 286
      - other patterns, 290
      - overview, 285
      - Pattern, 286–287
    - animation helpers, 292
    - Artistic filters
      - GIMPpressionist, 275–279
      - overview, 273–275
    - Blur filters, 257–259
    - Combine filters, 272–273
    - Distorts filters, 259–266
    - edge-detection filters, 270–271
    - Enhance filters, 259
    - Filters menu, 257
    - generic filters, 272
    - image window filters vs. Toolbox Xtns, 255
    - Light and Shadow filters
      - flares and sparkles, 266–267
      - lens and glass effects, 267–268
      - overview, 266
      - shadows, 267
    - Map filters
      - Bump Map filter, 279–280
      - Displace filter, 280–283
      - mapping toys, 283
      - overview, 279
      - tiling, 284
      - Warp filter, 284–285
    - noise filters, 268–269
    - overview, 255
    - tools vs. plug-ins, 256–257
    - using to make web pages, 290–292
  - Filters, Toolbox's Xtns menu, 294
  - Filters menu, 256, 418
  - Finding Contiguous Regions option, 17
  - Finding Similar Colors options, 173
  - fink method, 497
  - fire photoshop, 477
  - first project, 23–26
  - first startup, 507
  - Fixed aspect ratio, 44
  - Fixed Aspect Ratio option, 44
  - Flames filter, 286
  - Flare intensity, 267
  - FlareFX filter, 266
  - flares, 266–267
  - flat light, using Overlay for, 386–387
  - Flatland, Toolbox's Xtns menu, 294
  - Float feature, Select menu, 169
  - floating selection, 98
  - Focus Follows Mouse option, 497
  - folders, 23, 467, 475
  - FONT widget, 439
  - Force auto-hinter option, 80
  - Force baseline JPEG control, 38
  - foreground color (FG), 85, 140, 267, 427
  - Foreground Object Extractor, 213
  - foreground objects
    - extracting with SIOX, 213–216
    - highlighting, 196
  - Fractal Explorer, 288
  - Fractal Trace, GIMPpressionist, 283
  - fractals, 286, 288, 295
  - Frame disposal option, 114
  - Free filter, 260
  - Free Select (Lasso) tool, 138, 172
  - Free Select tool, 165, 192, 213, 267
  - free select, with Lasso tool, 138
  - FreeBSD, installing GIMP on, 500–501
  - Freehand button, 54
  - Frequently Asked Questions (FAQ)
    - documents, 477
  - From Path feature, Select menu, 169
  - From Selection button, 43
  - From selection option, 44
  - full tonal range, 47
  - fuzzy borders, 334–340
  - fuzzy brush, 233
  - fuzzy lines, drawing, 128–129
  - Fuzzy Select tool, 174, 469
- ## G
- g\_free(layers), 456
  - g\_malloc, 450
  - g\_new, 450
  - Gamma option, 464
  - gamma value, 48
  - Gaussian Blur filter, 196, 257, 352, 395
  - General Public License (GPL), 486
  - General tab, GIMPpressionist, 277

- generic filters, 272
- get\_foreground() function, Python, 440
- Gfig, 289
- ghostly images, 398
- GIF comment option, Save as GIF dialog, 41
- GIF Export dialog, 114
- GIMP
  - first project, 23–26
  - grouping dialogs by docking, 10–12
  - what it is, 1–2
  - windows
    - image windows, 5–10
    - layers dialog, 3–4
    - overview, 2
    - tool options dialog, 4
    - toolbox window, 2–3
- GIMP Announce mailing list, 478
- GIMP brush, 125, 369, 372
- Gimp C development environment,
  - installing, 421, 464
- GIMP Developer mailing list, 478
- GIMP Docs mailing list, 478
- gimp fire, 477
- Gimp FreeType plug-in, 490
- Gimp icon, Applications folder, 496
- GIMP Images mailing list, 478
- Gimp installation folder, 497
- gimp object, Python, 440
- GIMP OS2 mailing list, 478
- GIMP Perl mailing list, 478
- GIMP Procedure Browser (PDB), 442
- Gimp ToolKit (GTK), 471
- GIMP User mailing list, 478
- GIMP User's Group (GUG), 477, 482
- GIMP Web mailing list, 478
- Gimp-[version number].dmg file, 495
- Gimp-[version].dmg file, 496, 498
- gimp[version]setup folder, 489
- gimp[version]setup.zip file, 488
- gimp\_drawable\_get(layer\_id), 457
- gimp\_image\_get\_layers file, 456
- gimp-bucket-fill procedure, Perl, 442
- gimp-context-push, 433
- gimp-context-set-background line, 433
- gimp-drawable-fill procedure, Perl, 442
- gimp-edit-blend command, 435
- gimp-edit-bucket-fill procedure, Perl, 442
- gimp-edit-clear line, 433
- gimp-edit-fill line, 433
- gimp-edit-fill procedure, Perl, 442
- Gimp-Print position, 464
- gimp-python interpreters, 475
- gimprc
  - default brush, pattern, palette, gradient, and font, 469
  - default threshold, 469
  - Help browser, 470
  - help locales, 470
  - overview, 468
  - perfect mouse, 469
  - plug-in paths, 470–471
  - restoring keyboard shortcuts, 470
  - saving modified images, 470
  - share palette, 469
  - stingy memory use, 469
  - tear-off menus, 470
- GIMPressionist, 275–279
- gimp-splash.png file, 474
- gimptool --build autcropall.c, 454
- gimptool --install autcropall.c, 454
- gimptool program, 420
- GIMPwin Users mailing list, 478
- gimp-x.y.z.tar.bz2, 503
- gimp-x.y.z.tar.gz, 503
- glass effects, 267–268
- Glass Tile effect, 268
- Gnome project, 480
- GNU General Public License, 446, 489–490, 496
- GNU Image Manipulation Program. *See* GIMP
- GPL (General Public License), 486
- grab handles, 155
- gradients, 143–146, 208, 267, 377–378, 412, 435, 439, 481
- Grain extract mode, 347, 394
- Grain merge mode, 347, 403
- Grayscale mode, 40, 308
- grayscale or black and white, 306–311
  - Channel Mixer, 311
  - Decompose dialog, 309–310
  - Desaturate function, 308
  - Grayscale mode, 308
  - Hue-Saturation dialog, Colors menu, 308
  - methods of measuring brightness, 307
- Green color, mask thumbnail's border in the Layers dialog, 340
- Grid pattern, 286
- Grow feature, Select menu, 170
- Grow option, Interactive Warp filter, 263
- GTK (Gimp ToolKit), 471
- GTK libraries, 494
- GTK toolkit, 480
- gtk+[version].exe file, 489
- gtk+[version]setup.zip file, 488
- gtk-bookmarks.sav file, 474
- gtkrc, changing preview sizes with, 471–472
- GTK-Wimp, 489
- GUG (GIMP User's Group), 477, 482
- guides, 6

guillotine operation, 292  
Gutenprint print plug-in, 304

## H

handles, Bezier path, 184  
hard edges, 121–122, 135, 222  
Hard light mode, 346, 386–387  
hard-edged brush, 224  
hash (#) comment character, 439  
haze, using Burn to cut through, 387–388, 390  
Height field, 120  
Height option, 261, 462  
height parameter, 450  
Help browser, gimprc, 470  
help system, 16  
hidden files, 369  
highlight input slider, 48  
highlight slider, 49  
highlighting foreground objects, 196  
Highlights setting, 223  
Hinting option, 80  
Holdness parameter, 268  
horizontal gradient map, 282  
horizontal guide, 157  
horizontal mapping, 280  
hot spots, 83  
HSV (Hue, Saturation, and Value), 88, 302–303, 327–329  
HSV Noise, 268  
Hue control, 268  
Hue mode, 347  
Hue slider, 248  
Hue-Saturation dialog, Colors menu, 308  
Hurl filter, 269

## I

ICO (Windows Icon format), 36  
if clause, 440  
IFS Fractal filter, 286  
Illusion, GIMPpressionist, 283  
image brushes, vs. parametric brushes, 126  
image coordinates, vs. layer coordinates, 455–456  
image dimensions, 457  
image hoses, 374  
image modes, 39  
Image object, Python, 440  
Image Size options group, 19  
image thumbnails, 14  
Image Type field, 419  
IMAGE widget, 439  
image window filters, vs. Toolbox Xtens, 255  
image windows, 6, 8–10, 20–21  
    context menus and tear-offs, 7–8  
    image window menus, 8–10  
    overview, 5–7  
ImageMap, 290  
Image/Output Settings, 461–464  
images. *See also* filters; photos  
Image-sized layers, 163  
importing animated brushes, 374  
Include window decoration controls, 466  
Incremental option, 127, 373  
Indent option, 81  
indentation, Python, 440  
Independent RGB filter, 269  
indexed color, 317–322  
    converting to indexed mode, 318–319  
    dithering an indexed image, 319–320  
    overview, 317–318  
    redesigning for better indexed results, 320–322  
Indexed Colors dialog, 40  
indexed format, 34  
indexed mode, 40  
indexed palette, 301  
indexed PNG, 39–41  
initial zoom ratio, 20  
Ink Pen Tool, 131–133  
input devices, 22, 133  
Input Levels, Brightness-Contrast tool, 47  
input sliders, 48  
installer (.exe) files, 491  
installing GIMP  
    on Linux or Unix, 499  
        Debian variants, 500  
        FreeBSD and NetBSD, 500–501  
        Linspire, 500  
        Mandriva, 500  
        other Unix versions, 501  
        overview, 499  
        Red Hat, 500  
        SuSE, 500  
    on Mac, 493–498  
    on Windows, 487–496  
Intelligent Scissors, 186–187  
Interactive Warp filter, 261  
Interface category, 470  
interface widgets, 437, 444  
Interfaces category, 471  
Interlace option, Save as GIF dialog, 40  
Interlacing option, Save as PNG dialog, 41  
internationalization, 436, 454  
Interpolation option, Rotate tool, 61  
Interpolation Quality controls, 30  
Intersect mode, 191  
Inverse, 267  
Invert bumpmap option, Bump Map, 280  
Invert feature, Select menu, 169  
inverting colors, 56

Iterations, GIMPpressionist, 285  
IWarp filter, 261, 415

## J

Jigsaw pattern, 287  
Join style, 137  
JPEG (Joint Photographic Experts Group), 34, 37–39

## K

Keep aspect tool options, 32  
Keep transparency toggle button, 152  
Keep transparent button, 102  
kestrel, 64  
key bindings, 473  
keyboard shortcuts, 14–15, 470  
killcols array, 450  
killrows array, 450, 452

## L

Land, Toolbox's Xtns menu, 294  
landscape image, 19  
Landscape option, 462  
Laplace filter, 270  
Lasso (Free Select) tool, 138, 172  
Launch The GIMP, 490  
Lava pattern, 290  
layer boundary, 240  
layer coordinates, vs. image coordinates, 455–456  
Layer Fill Type, 120, 354  
layer masks, 103, 203–212, 390  
Layer menu, 9  
layer modes, 315  
  Addition mode, 344–345  
  Burn mode, 345–346  
  Color mode, 347  
  creating depth with, 351–356  
  Difference mode, 344–345  
  Divide mode, 345  
  Dodge mode, 345–346  
  Grain extract mode, 347  
  Grain merge mode, 347  
  Hard light mode, 346  
  Hue mode, 347  
  Lighten only mode, 347  
  Multiply mode, 345  
  Overlay mode, 345–346  
  overview, 341–344  
  Saturation mode, 347  
  Screen mode, 345–346  
  self-compositing with  
    overview, 384–385  
    using Overlay or Burn to cut through haze, 387–390

  using Overlay or Hard light when light is flat, 386–387  
  using Screen mode for dark images, 385–386

  Soft light mode, 346  
  Subtract mode, 344–345  
  Value mode, 347

Layer Name field, 120

Layer object, Python, 440

Layer Offset

  black velvet paintings: Difference mode, 393–394

  colored pencil drawings: Divide mode, 391–393

  embossing: Grain extract, 394

  overview, 390–391

Layer preview button, 106

layer stack, 103

Layer to Imagesize option, 163

LAYER widget, 439

layers. *See also* Layers dialog

  adding patterns to

    Clouds, 285–286

    Fractal Explorer, 288

    Gfig, 289

    Nature, 286

    other patterns, 290

    overview, 285

    Pattern, 286–287

  drawing using, 119–120

  loading all images as, 405

  making simple GIF animations (project)

    testing animation, 113

    for Web, 113–115

  overview, 73–75

  performing operations on, 94–95

  simple effects using, 89–93

  using for copy and paste, 96–100

  using Move tool, 82–85

  using to add text, 76–81

Layers dialog

  buttons, 104

  context menus

    Add Alpha Channel, 106

    Alpha to Selection, 107

    Edit Layer Attributes, 106

    Flatten option, 107

    Layer Boundary Size, 106

    layer controls, 106

    Layer to Image Size, 106

    Mask options, 106

    Merge option, 107

    overview, 105

    Scale Layer, 106

    Text Tool item, 105

- Layer mode, 102
  - Layers list, 103–104
  - Opacity slider, 102–103
  - overview, 101
  - tabs, 102
  - title area, 102
  - ldconfig command, 506
  - lens effects, 267–268
  - lens warp, 410
  - Letter spacing option, 81
  - library list, 504
  - LIC (Van Gogh) filter, 274
  - LICENSE.txt, 496
  - light, increasing by additive stacking, 406
  - Light and Shadow filters
    - flares and sparkles, 266–267
    - lens and glass effects, 267–268
    - overview, 266
    - shadows, 267
  - Lighten only mode, 347
  - lighter effect, using Screen for, 399
  - Lighting (degrees) slider, 432
  - Lighting Effects, 267
  - lightness of a pixel, 307
  - Lightness slider, 70
  - Limit line width option, Engrave filter, 261
  - Line art, 463
  - Line Nova pattern, 290
  - Line spacing option, 81
  - Line Style, 137
  - Line Width, 137
  - Linear button, 49, 54
  - Linear filter, 258
  - line-painting tool, 136
  - lines and curves
    - Airbrush tool, 130
    - brushes, 122–126
    - drawing fuzzy or smooth lines, 128–129
    - drawing tool options, 127–128
    - hard edged lines, 121–122
    - Ink Pen Tool, 131–133
    - overview, 121
    - Paintbrush tool, 128–129
    - Pencil tool, 121–122
  - Linspire, installing GIMP on, 500
  - Linux distribution, 499
  - Lisp, 422, 427, 431
  - lists, 431
  - Logarithmic button, 49, 54
  - Logos, Toolbox's Xtns menu, 294
  - long exposures, 247
  - Loop forever option, 114
  - loops, 457
  - lossy compression, 34
  - Lower filter, 259
  - lower parameter, 437
  - LTR, Text Editor window, 79
  - luminosity, 311
  - Luminosity threshold, 267
- ## M
- Mac, installing GIMP on, 493–498
  - Magic Wand tool, 174
  - magnitude map, GIMPpressionist, 285
  - mailing list page, 478
  - main() function, 446
  - Main Keyboard tab, 22
  - make install phase of build, 506
  - make phase of build, 505
  - Make Seamless, 284, 376
  - Mandriva, installing GIMP on, 500
  - Manually chain link icon, 21
  - Map filters
    - Bump Map filter, 279–280
    - Displace filter, 280–283
    - mapping toys, 283
    - overview, 279
    - tiling, 284
    - Warp filter, 284–285
  - Map Object, GIMPpressionist, 283
  - mapping toys, 283
  - Marble Madness preset, GIMPpressionist, 277
  - marching ants, 168
  - masks, 103, 106
    - fuzzy borders, 334–340
    - interpreting mask border color, 340
    - overview, 333
    - text fade out, 333–334
  - Max depth option, Interactive Warp filter, 263
  - Maze pattern, 287
  - Measure tool, 405
  - Media size option, 461
  - Media type option, 461
  - memory, 485
  - menu bar, image window, 6
  - menu entries, 473
  - menu mnemonics, 436
  - menurc file, 473
  - menus
    - context menus, 7–8
    - image window menus, 8–10
  - Merge Down option, 106, 110, 351
  - Merge option, 107
  - metering, 411
  - middle input level slider, 48
  - midtone input slider, 48
  - Midtones, 250
  - Mirror filter, 259
  - mirror radio box, 444

- mirror sites, 488
- Mirrorsplit, Perl, 443
- Miter limit, 137
- Mode: Disabled option, 133
- Mode menu, 341, 414
- Mode option, 127, 133, 145
- Mode selections, 140
- Mode tool options, 146
- Modes, 188
- modifications, self-compositing with
  - adding blurs and other tricks for artistic effects, 395
  - Layer Offset, 390–394
  - overview, 390
- monochrome images and sepia photos, 311–315
  - automatic conversion with Old Photo filter, 312–313
  - manual conversion for fine control, 313–315
- Mosaic filter, 263
- Motion Blur filter, 257, 395
- Mouse Cursors preference, 20
- Move mode, 185
- Move option, Interactive Warp filter, 262
- Move the current layer option, Move tool, 83
- Move tool, 25, 82–85, 155, 238, 391, 406, 413
- multiplicative stacking, 407–408
- Multiply layer mode, 280
- Multiply mode, 345, 430
- multipoint perspective
  - overview, 360
  - single-point perspective, 361
  - two-point perspective, 361–363
- my-sphere script, 436

## N

- NASA (National Aeronautics and Space Administration), 482
- National Oceanic and Atmospheric Administration (NOAA), 482
- National Park Service website, 482
- Nature pattern, 286
- navigation control, 7
- Navigation preview size controls, 471
- Neon filter, 270
- NetBSD, installing GIMP on, 500–501
- New brush, 125
- New Image dialog, 119
- New Image option, 108
- New Layer button, 99, 108, 204, 382
- New layer button, Layers dialog, 104, 119
- New Printer option, 461
- Newsprint filter, 263
- Next button, 489
- NL Filter, 259
- NOAA (National Oceanic and Atmospheric Administration), 482
- noise, reducing, 405
- noise filters, 268–269
- Non-aligned mode, 237
- None feature, Select menu, 169
- None mode, drawing tool, 146
- non-lossy raw format, 34
- Normal layer mode, 342
- normal mode, 74, 146
- Normalize function, Auto menu, 57

## O

- obtaining GIMP, 486
- Offset fields, 162
- Offset mode, drawing tool, 145
- Offset slider, 435
- Oilify filter, 274
- Old Photo filter, 312–313
- One frame per layer option, 114
- Only background filter, 265
- Only foreground filter, 265
- opacity, 90, 102–103, 127, 140, 204, 267, 384–385, 409, 435
- opaque background layer, 135
- Open brush as image option, 125
- Open Location option, 28
- Open Pattern as Image option, 143
- Open Recent menu option, 28
- Open Recent submenu, 8
- open source software, 1, 424, 504
- opening files, 27–28
- Optimize method, 38
- Optimize option, 292
- OPTION widget, 438
- orientation, setting when printing, 462–463
- Orientation preset, GIMPpressionist, 277
- OS X program, 498
- outlining selections, 136–137
- Output Levels, Brightness-Contrast tool, 47
- Output type, 463
- Overlay, 345–346, 388
  - using for dark images, 397–398
  - using when light is flat, 386–387

## P

- Pagecurl filter, 263
- Paint Options Shared Between Tools, 17
- Paintbrush tool, 128–129, 152, 232, 414
- palettes, 39, 439, 469
- panning around images, 7



- panoramas, stitching
  - adding rest of images, 415
  - adjust layer mask, 414
  - adjusting layer mask, 414
  - adjusting with other tools, 415
  - calculating expected new image size, 411
  - deciding on resolution, 411
  - final adjustments, 415
  - fine-tuning position of second image, 413
  - loading first two images, 412
  - making Gradient layer mask, 412
  - overview, 409–410
  - shooting images, 410–411
- paper type, setting when printing, 461
- parameters, 432
- Parameters tab, Fractal Explorer, 288
- parametric brushes, 125–126
- parasiterc file, 474
- parenthesis-delimited lists, 468
- partial transparency, 35
- Path to Selection option, 178
- paths
  - adding nodes or segments and moving paths, 185
  - curved, 184–185
  - defining, 175–177
  - moving or modifying an existing path, 185–186
  - Paths dialog, 177–182
- patterns, 142–143, 438
  - adding to layers
    - Clouds, 285–286
    - Fractal Explorer, 288
    - Gfig, 289
    - Nature, 286
    - other patterns, 290
    - overview, 285
    - Pattern, 286–287
  - making, 374–376
- Patterns, Toolbox's Xtns menu, 294
- Patterns dialog, 153
- PCD (proprietary Photo CD format), 36
- PDB (GIMP Procedure Browser), 442
- PDF vector graphics formats, 36
- Pencil line, 129
- Pencil tool, 121–122, 128
- Percent option, 462
- perfect mouse setting, gimprc, 469
- Perl, writing GIMP scripts with, 423, 465
- Perl script, 443–444
- Perlotine filter, 291
- Personal GIMP Folder screen, 507
- Perspective filter, 267
- Perspective tool, 357
- perspective transformation, 153–161
- PGM (Portable GrayMap) format, 277
- Photocopy filter, 274
- Photograph option, 463
- photos
  - brightening and darkening
    - brightness-contrast, 44–47
    - curves, 52–54
    - levels, 47–52
    - other exposure adjustments, 54, 56–57
  - overview, 44
  - cropping, 42–44
  - GIF and indexed PNG quality settings, 39–41
  - image file types
    - GIF, 34
    - JPEG, 34
    - other formats, 36
    - overview, 33
    - PNG, 35
    - TIFF, 35
    - XCF, 35
  - JPEG quality settings, 37–39
  - opening files, 27–28
  - overview, 27
  - red-eye, 66–71
  - rescaling
    - deciding on size, 30–32
    - overview, 29–30
    - scale tool, 32
  - rotating
    - free rotation, 58–61
    - by multiples of 90 degrees, 58
    - overview, 58
  - saving files, 32–33
  - sharpening, 62–66
- Pick a layer or guide option, Move tool, 83
- Pick Colors from the Image option, 108
- Pick filter, 269
- Pinch filter, 265
- pipes, 372
- pixel regions
  - using, 450–451
  - writing to, 453–454
- pixel resolutions, 18
- Pixelize filter, 258
- pixels, 192
- Pixels per Inch (PPI), 462
- Placement control preset, GIMPpressionist, 277
- Plasma filter, 286
- Playback option, 292
- Plug-in Browser, 3, 418–419, 461, 477
- plug-in paths, gimprc, 470–471
- pluginrc file, 474

- plug-ins, 376. *See also* C plug-in, writing; scripts
    - finding
      - overview, 419, 461
      - Plug-In Registry, 419
      - using web search, 420
    - installing
      - C plug-in installation, 420–421, 463
      - Gimp C development environment installation, 421, 464
      - overview, 420
      - script installation, 420
      - overview, 418, 460
      - Plug-In Browser, 418–419, 461
    - plug-in-zealouscrop, 446
    - PNG, 35, 39–41
    - Point lighting type, 267
    - Polar Coords filter, 264
    - Polar mode, GIMPpressionist, 283
    - Polygonal box, Paths tool options, 185
    - Portable GrayMap (PGM) format, 277
    - Portable PixMap (PPM) format, 277
    - portrait image, 19
    - Portrait option, 462
    - portrait photography, 219
    - position of images, setting when printing, 462–463
    - Position section, 461
    - Posterize filter, 274
    - posterize images, 55
    - PostScript Level 2 setting, 461
    - PostScript vector graphics formats, 36
    - PPI (Pixels per Inch), 462
    - PPM (Portable PixMap) format, 277
    - Predator filter, 274
    - Preference dialog, 467
    - Preferences dialog, 30
    - Preferences window, 15, 468
    - prefix notation, 431
    - Presets tab, GIMPpressionist, 278
    - Pressure option, Airbrush, 130
    - Pressure sensitivity options, 222
    - preview, Layer, 103
    - Preview button, 40, 246
    - Preview checkbox, 257
    - Preview option, 45
    - Preview Size item, 323
    - Preview Size options, 471
    - preview sizes, changing with gtkrc, 471–472
    - Print Color Adjust dialog, 463
    - Printer name menu, 464
    - printing
      - to file, 464
      - Image/Output Settings, 463–464
      - making print, 463
      - overview, 459–460
      - setting image's position, orientation, and size, 462–463
      - setting print size, paper type, and resolution, 461
      - setting printer model, 460–461
    - printrc file, 474
    - Procedural DataBase, 441
    - Procedure Browser, 442
    - programming interfaces, GIMP, 486
    - progress-bar, 507
    - Progressive method, 38
    - proprietary Photo CD format (PCD), 36
    - PSD (proprietary Photoshop format), 36
    - Py-Slice filter, 291
    - Python, writing GIMP scripts with, 422–423, 465
    - Python module, 476
    - Python script, 439–441, 483
    - python\_sphere routine, 441
- Q**
- Qbist pattern, 287
  - Quality setting, 36
  - Quick Start Version, 486
  - QuickMask, 6, 191–195, 326–327
  - quotation marks, 426
- R**
- Radial filter, 258
  - RADIO, 437
  - Radius filter, 259, 266
  - Radius parameter, 64, 259
  - Raise and Lower buttons, Layers dialog, 104
  - Raise button, 111
  - Random Hue, 267
  - Random option, 267, 373
  - Random Saturation, 267
  - Randomization filter, 269
  - range selection sliders, 377
  - Ranks entry controls, 373
  - Rate option
    - Airbrush tool, 130
    - Smudge tool, 228
  - Rate slider, Convolve tool, 241
  - raw proprietary formats, 36
  - README.txt file, 496
  - Ready To Install dialog, 489
  - realism and multipoint perspective
    - overview, 360
    - single-point perspective, 361
    - two-point perspective, 361–363
  - Rectangle Select tool, 135, 153, 189
  - Recursive filter, 259
  - Red, Green, and Blue (RGB), 88
  - Red color, mask thumbnail's border in the Layers dialog, 340

- Red Hat Linux, installing GIMP on, 500
  - red pixel, 342
  - reduced tonal range, 47
  - Reflective filter, 265
  - Refresh brushes, 126, 370
  - Refresh button, 375
  - Refresh Scripts, 429, 454
  - regions
    - defining, 135–136
    - filling, 138–139
  - register function, 441, 444
  - Registered mode, 237
  - registering images, using Difference mode, 405–406
  - registering script, 435–437, 478
  - Remove button, 28
  - Remove option, Interactive Warp filter, 262
  - Render Map, Toolbox's Xtns menu, 294
  - Repeat button, 257
  - Repeat filter, 269
  - Repeat setting, 435
  - Replace default mode, 188
  - requirements for GIMP, 485–486
  - rescaling photos
    - deciding on size, 30–32
    - overview, 29–30
    - scale tool, 32
  - Reset all Filters button, 257
  - Reset button, 45
  - Reset filter, 259
  - Reset to default values button, 257
  - Re-Show button, 257
  - Resize Layers menu, 163
  - Resize Layers option, 162
  - Resize option, 44
  - resizing, 29
  - resolution, 19
    - increasing by averaging, 409
    - for panoramas, 411
    - setting when printing, 461
  - Resolution section, Scale Image dialog, 30
  - Resource Consumption preferences, 13
  - resources, additional
    - art sites and imaging contests, 482–483
    - finding source code at cvs.gnome.org, 481
    - mailing lists, 478
    - official GIMP documentation
      - developer documentation, 476
      - other official GIMP user documentation, 476
    - overview, 475
    - other tutorial sites, 477–478
    - overview, 475
    - reporting bugs, 479–481
    - where to find freely available images, 481–482
    - Wiki, 477
  - Resynthesizer plug-in, 376
  - Reverse mode, drawing tool, 145–146
  - RGB (FG to BG), 145
  - RGB (Red, Green, and Blue), 88
  - RGB and CMY colors
    - additive colors, 298–299
    - color depth, 300–301
    - indexed palette, 301
    - representing RGB color in web pages, 301–302
    - subtractive colors, 299
  - RGB Noise filter, 269
  - Ripple filter, 264
  - Rippling option, 292
  - RLE method, 246
  - Rotate filter, 260
  - Rotate tool, 94, 410
  - rotating text layer, 94
  - Rotation angle, GIMPpressionist, 285
  - Rounded Rectangle feature, Select menu, 170
  - RTL, Text Editor window, 79
  - run procedure, Zealous Crop, 451
- S**
- Sample merged option, 142
  - SANE (Scanner Access Now Easy), 465
  - saturation, 69, 268, 328, 347
  - Save a Copy option, 33
  - Save as Brush Pipe dialog, 372
  - Save As dialog, 85
  - Save As option, 33
  - Save background color option, Save as PNG dialog, 41
  - Save color values from transparent pixels option, Save as PNG dialog, 41
  - Save comment option, Save as PNG dialog, 41
  - Save creation time option, Save as PNG dialog, 41
  - Save EXIF data option, 38
  - Save gamma option, Save as PNG dialog, 41
  - Save in folder popup, 33
  - Save layer offset option, Save as PNG dialog, 41
  - Save resolution option, Save as PNG dialog, 41
  - Save Settings button, 22
  - Save thumbnail option, 39
  - Save tool options now, Tool Options category, 469
  - Save XMP option, 39
  - saving
    - files, 32–33
    - modified images, gimprc, 470
    - selections using channel masks, 197–199, 201

- Sawtooth wave mode, drawing tool, 146
- Scale button, 30
- Scale Image dialog, 29
- Scale Layer, 106
- scale menu, 6, 31
- Scale option, 161
- Scale tool, 32, 99
- scaling images, 29. *See also* rescaling photos
- Scaling option, 17
- Scanner Access Now Easy (SANE), 465
- scanner drivers, 465
- scanning from GIMP, 465
- .scm (scheme) extension, 424
- Screen mode, 345–346, 385–386, 399
- screen shots, 465–466
- Screenshot option, 465
- script author, 436
- Script Progress option, 90
- Script-Fu language, 256, 422, 465
- Script-Fu module, 476
- script-fu-menu-register, 425
- script-fu-my-sphere, 426
- script-fu-register command, 426, 430, 436
- script-fu-sphere function, 426, 430, 432–435, 473–476
- scripts, 256. *See also* plug-ins; scripts, writing
  - Perl script, 443–444
  - Python script, 439–441, 483
  - user interface options for
    - ADJUSTMENT, 437–438
    - BRUSH, 438
    - COLOR, 438
    - DIRNAME, 438
    - ENUM, 438
    - FILENAME, 438
    - FONT, 439
    - GRADIENT, 439
    - IMAGE, DRAWABLE, LAYER, 439
    - OPTION, 438
    - overview, 437, 478
    - PALETTE, 439
    - PATTERN, 438
    - RADIO, 437
    - TEXT, 437
    - TOGGLE, 437
    - VALUE, 437
- scripts, writing
  - overview, 422
  - with Perl, 423, 465
  - with Python, 422–423, 465
  - reading and modifying script
    - bringing shadow back, 430
    - editing copy of script, 424–425, 467
    - finding original script, 424, 466
    - GIMP peculiarities, 428–430, 472
    - making change, 427, 469–470
    - making script your own, 425–427, 469
    - overview, 423–424
    - Sphere script, 431–437, 473–478
    - trying your fix, 428, 470
  - with Script-Fu language, 422, 465
- Search button, 495
- Search Downloads feature, 495
- Search page, 479
- Seascape option, 462
- segments, 377
- Select by Color, 171–173
- Select Destination dialog, 489–490
- Select File Type option, 33
- selection boundary, 238
- Selection Editor feature, Select menu, 169
- Selection to Path option, 178
- selections
  - Bezier Paths
    - adding nodes or segments and moving paths, 185
    - curved paths, 184–185
    - defining a path, 175–177
    - moving or modifying an existing path, 185–186
    - overview, 174–175
    - Paths dialog, 177–182
  - extracting foreground objects with SIOX, 213–216
  - Fuzzy select, 174
  - highlighting foreground objects, 196
  - Intelligent Scissors, 186–187
  - layer masks, 203–212
  - modifying with selection modes, 187–191
  - outlining, 136–137
  - overview, 167
  - QuickMask, 191–195
  - saving using channel masks, 197–201
  - Select by Color, 171–173
  - working with
    - marching ants, 168
    - moving selections, 170–171
    - overview, 168
    - Select menu, 168–170
- Selective Gaussian Blur filter, 259
- self-compositing
  - with layer modes
    - overview, 384–385
    - using Overlay or Burn to cut through haze, 387–390
    - using Overlay or Hard light when light is flat, 386–387
    - using Screen mode for dark images, 385–386

- with modifications
  - adding blurs and other tricks for artistic effects, 395
  - Layer Offset, 390–394
  - overview, 390
- Semi-Flatten filter, 292
- Sensitivity option, Ink Pen, 132
- sepia photos. *See* monochrome images and sepia photos
- sessionrc, 472
- settings, customizable
  - color management, 21
  - default grid, 20
  - default image, 18–19
  - display, 21
  - environment, 13–14
  - folders, 23
  - help system, 16
  - image windows, 20–21
  - input devices, 22
  - overview, 13
  - theme, 15–16
  - tool options, 16–18
  - toolbox, 18
  - user interface and keyboard shortcuts, 14–15
  - window management, 22
- Setup Printer option, 461
- Severity field, 480
- SGI, 501
- shadow input slider, 48
- shadow parameter, 450
- shadow slider controls, 49
- shadow width, 434
- shadows, 267. *See also* Light and Shadow filters
  - overview, 357–358
  - transparency, 359–360
  - using object to cast own shadow, 358–359
- Shape mode, drawing tool, 145
- Shape option, 132, 435
- shapes, drawing
  - Bucket Fill tool, 139–142
  - defining regions, 135–136
  - filling regions, 138–139
  - free select with Lasso tool, 138
  - Gradients tool, 143–146
  - outlining selections, 136–137
  - overview, 135
  - patterns, 142–143
- sharing palettes, gimprc, 469
- Sharpen feature, Select menu, 169
- Sharpen filter, 259
- Sharpen Mask, 62
- Shift filter, 265
- Shift-clicking, 185
- Shift-drag, 185
- Show Advanced Options button, 480
- Show brush outline option, 122
- Show Image Selection option, 102
- Show interactive boundaries option, 185, 187
- Show Layer Mask, Layers dialog, 210
- Show Package Contents, 425
- Show paint tool cursor, 122
- Show Preview in image window checkbox, 37
- Shrink feature, Select menu, 170
- Shrink option, Interactive Warp filter, 263
- Simple Interactive Object Extraction (SIOX), extracting foreground objects, 213–216
- single open bracket ([ ]), 444
- single quote, 433
- single-point perspective, 361
- Sinus pattern, 287
- SIOX (Simple Interactive Object Extraction), extracting foreground objects, 213–216
- size. *See* rescaling photos
- size of images, setting when printing, 462–463
- Size preset, GIMPpressionist, 277
- Size slider, 462
- Size tab preset, GIMPpressionist, 277
- Slur filter, 269
- Small Tiles, GIMPpressionist, 284
- Smear filter, 265
- smoothing, 38, 128–129, 216, 260
- smudging, 226–227
- Snap to guides, Preferences, 158
- Sobel filter, 270
- Soft light mode, 346, 396
- Softglow filter, 274
- Solid colors, 463
- Solid Noise filter, 286
- source, building from
  - configure phase, 504–505
  - getting source, 503–504
  - make install phase, 506
  - make phase, 505
  - overview, 503
- Source option, 236
- source pixel region, 453
- source region, 234
- spacing, 125, 369, 373
- sparkles, 266–267
- Sphere Designer, 290, 294
- Sphere script
  - overview, 431, 473
  - registering script, 435–437, 478
  - script-fu-sphere function, 432–435, 473–476

- Sphere scripts, 425
  - Spike options, 267
  - Spinning Globe option, 292
  - Spiroimp pattern, 290
  - splashes folder, 475
  - Spread filter, 269
  - square brackets ([ ]), 441
  - stable version, 475
  - stacking images
    - additive stacking, 406
    - averaging, 409
    - loading all images as layers, 405
    - multiplicative stacking, 407–408
    - overview, 404–405
    - reducing noise, 405
    - registering images using Difference mode, 405–406
  - stamps, 370
  - Start Menu folder, 490
  - Status options, 480
  - Stencil Carve, 294
  - Stencil Chrome, 294
  - Step size, GIMPpressionist, 285
  - stitching, 410
  - Stock.XCHNG, 482
  - Strength filter, 266
  - Stretch Contrast function, Auto menu, 57
  - Stretch HSV function, Auto menu, 57
  - stroking, 136–138, 178
  - Subsampling control, 38
  - Subtract layer mode, 342
  - Subtract mode, 191, 344–345, 394, 400–402
  - subtractive colors, 299
  - Super GIMP Pencil, 122
  - Supersampling option, Rotate tool, 61
  - SuSE, 500
  - SVG format, 36
  - Swap File, 507
  - Swap filter, 259
  - Swirl option, Interactive Warp filter, 263
  - Swirl-Tile, Toolbox's Xtns menu, 295
  - Swirly, Toolbox's Xtns menu, 295
  - synaptic graphical program, 500
  - system directory, 506
  - system-wide interpreters folder, 474
- T**
- Tab menu, 11
  - tabs, Layers dialog, 102
  - Tagged Image File Format (TIFF), 35
  - tarball, 503
  - tear-off menus, gimprc, 470
  - tear-offs, 7–8
  - Technology Without An Important Name (TWAIN), 465
  - Template drop-down menu, 117
  - templaterc file, 474
  - Templates dialog, 18
  - temporary variables, 434
  - terminal programs, 504
  - Terminal.app, 497
  - terminals, 504
  - Text along Path button, 293
  - Text along path option, 81, 87
  - text editor, 425
  - Text Editor window, 78
  - text fade out, 333–334
  - text layer, 75, 94
  - text tool, 78–81, 86, 105, 293
  - TEXT widget, 437
  - texture, 403
  - theme, 15–16
  - themerc file, 474
  - 3D Truchet, Toolbox's Xtns menu, 294
  - Threshold dialog, 54
  - Threshold option, Interactive Warp filter, 263
  - Threshold parameter, Unsharp Mask, 64
  - Threshold slider, 140
  - Threshold tool, 315–317, 326–327
  - thumbnails, 14
  - TIFF (Tagged Image File Format), 35
  - Tigert Labs, 477
  - tiling, 143, 259, 280, 284, 375, 449, 451–452, 507
  - title area, Layers dialog, 102
  - tmp folder, 475
  - To Path feature, Select menu, 170
  - toggle selection boundaries, 168
  - Toggle Selection option, 146
  - TOGGLE widget, 437
  - tonal range, 50
  - tool buttons, 2
  - tool icons, 507
  - tool options, 16–18, 122
  - Tool Options category, 469
  - tool options dialog, 4, 11, 256
  - toolbox, 18
  - toolbox window, 2–3, 507
  - Toolbox Xtns, vs. image window filters, 255
  - ToolKit, GIMP, 486
  - tool-options folder, 475
  - Tools menu, 3, 9
  - tooltips, 3
  - touch ups
    - blurring backgrounds with Gaussian Blur, 245–247
    - blurring with Convolve tool, 242–244
    - Burn tool, 222–224
    - copying small regions, 237–240



- Curves tool, 252–253
  - Dodge tool, 221–222
  - fine-tuning clone job, 235–236
  - hue-saturation, 248–250
  - Levels tool, 251–252
  - options, 236–237
  - overview, 219–220
  - setting clone source, 232–234
  - sharpening with Convolve tool, 242
  - smudging blemishes away, 225–229
  - Transform Direction option, Rotate tool, 61
  - Transform layer, 97
  - Transform selection tab, 364
  - transparency, 34, 74, 89, 258, 359–360
  - Tree View tab, 418
  - triangle color selector, 305
  - Triangle wave mode, drawing tool, 146
  - Truchet, Toolbox's Xtns menu, 294
  - tubes, 374
  - Turn On X11 Focus Follows Mouse option, 497
  - tutorials, 478
  - Tux layer, 358
  - TWAIN (Technology Without An Important Name), 465
  - TWAIN Acquire menu, 465
  - two-point perspective, 361–363
  - Type option, Ink Pen, 132
- U**
- underscore, 426
  - Undo command, 30, 177
  - Undo History option, 30
  - Undo stack, 433
  - unitrc file, 474
  - units menu, 19
  - UnOptimize option, 292
  - unrelated images, compositing
    - making eerie colors with Burn, 402–403
    - overview, 396
    - using Addition to complement light and dark, 399–400
    - using Grain merge to add texture, 403
    - using Overlay for dark images, 397–398
    - using Screen to get lighter effect, 399
    - using Soft light for combining images, 396
    - using Subtract to make cut-out mask, 400–402
  - Unsharp Mask, 62, 240, 259
  - unzip utility, 489
  - Update option, GIMPpressionist, 276–277
  - Upside down option, 462
  - Use color from gradient option, 127
  - Use dynamic keyboard shortcuts option, 15
  - Use restart markers control, 38
  - user interface, 14–15, 471, 486
  - Utility Window, 22
- V**
- VALUE, STRING widget, 437
  - Value mode, 347
  - Value Propagate filter, 265
  - Van Gogh (LIC) filter, 274
  - vanishing point, 361
  - variables, in Script-Fu, 432
  - vector-drawing programs, 135
  - vertical lens gradient, 282
  - vertical mapping, 280
  - Vertical numbers, 246
  - verticals, 58
  - Video filter, 265
  - Video menu, 256
  - View menu, image window, 8
  - Vine brush, 124
  - Visibility icon, 103
  - visible pixel, 457
- W**
- Warp filter, 284–285
  - Watercolor selector, 306
  - Waterlevel option, Bump Map, 280
  - Waves filter, 264–265
  - Waves option, 292
  - Weave filter, 274
  - web pages
    - filters to help make, 290–292
    - representing RGB color in, 301–302
  - web-optimized palette, 39
  - Whirl angle slider, 265
  - Whirl filter, 265
  - White Balance function, Auto menu, 57
  - White color, mask thumbnail's border in the Layers dialog, 340
  - White level control, 259
  - white pixel, 342
  - Width field, 120
  - Width option, 462
  - Wikipedia, 477
  - Wilber icon, 498
  - wildcard character (\*), 419
  - Wind filter, 266
  - window management, 22
  - Window Manager Hints controls, 22
  - windows
    - image windows, 7–10
    - layers dialog, 3–4
    - overview, 2
    - positions, 472
    - tool options dialog, 4
    - toolbox window, 2–3

Windows, installing GIMP on, 487–491  
Windows Bitmap format (BMP), 36  
Windows Icon format (ICO), 36  
Work on copy filter, 260  
Writable folder, 368

**X**

X icon, 495  
X offset option, Bump Map, 280  
X Window System, 496  
X11User.dmg, 496, 498  
Xach Beane, 268  
Xach effect, 267  
XCF format, 35

XDarwin, 497  
XFree86, 497  
XSane menu, 465  
Xtns menu, 3, 294–295, 418

**Y**

Y offset option, Bump Map, 280  
YaST software screen, SuSE, 500

**Z**

Zealous Crop plug-in, 445–447, 451, 454,  
479–480  
zooming, 6–7, 20, 31, 176, 223, 258, 288