


SECOND EDITION

R E A L  W O R L D

Image Sharpening

with Adobe® Photoshop®, Camera Raw, and Lightroom®

INDUSTRIAL-STRENGTH PRODUCTION TECHNIQUES



Nuts and bolts of
sharpening

Creative-sharpening
techniques

Workflow strategies

BRUCE FRASER and JEFF SCHEWE

REAL WORLD IMAGE SHARPENING

with Adobe Photoshop, Camera Raw, and Lightroom,
Second Edition

BRUCE FRASER
JEFF SCHEWE



PEACHPIT PRESS
BERKELEY, CALIFORNIA

REAL WORLD IMAGE SHARPENING WITH ADOBE PHOTOSHOP, CAMERA RAW, AND LIGHTROOM, SECOND EDITION

Bruce Fraser and Jeff Schewe

Peachpit Press
1249 Eighth Street
Berkeley, CA 94710
510/524-2178
510/524-2221 (fax)

Find us on the Web at: www.peachpit.com

To report errors, please send a note to: errata@peachpit.com

Peachpit Press is a division of Pearson Education.

Published in association with Adobe Press

For the latest on Adobe Press books, go to: www.adobepress.com

Copyright © 2010 by Jeff Schewe

Editor: Rebecca Gulick
Production Editor: Lisa Brazieal
Copy Editor: Emily K. Wolman
Proofreader: Elle Yoko Suzuki
Compositor: WolfsonDesign
Indexer: Valerie Haynes Perry
Cover Design: Charlene Will
Cover Illustration: John Weber
Cover Photo: Jeff Schewe

Notice of Rights

All rights reserved. No part of this book may be reproduced or transmitted in any form by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. For information on getting permission for reprints and excerpts, contact permissions@peachpit.com.

Notice of Liability

The information in this book is distributed on an “As Is” basis, without warranty. While every precaution has been taken in the preparation of the book, neither the authors nor Peachpit Press 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 instructions contained in this book or by the computer software and hardware products described in it.

Trademarks

“Adobe,” “Adobe Bridge,” “Photoshop Camera Raw,” “Lightroom,” and “Photoshop” are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Peachpit was aware of a trademark claim, the designations appear as requested by the owner of the trademark. All other product names and services identified throughout this book are used in editorial fashion only and for the benefit of such companies with no intention of infringement of the trademark. No such use, or the use of any trade name, is intended to convey endorsement or other affiliation with this book.

13-digit ISBN: 978-0-321-63755-0

10-digit ISBN: 0-321-63755-0

9 8 7 6 5 4 3 2 1

Printed and bound in the United States of America

TABLE OF CONTENTS

Preface: Real World Sharpening.....vii

The Sharpening Workflow	viii
Sharpening and the Display.....	ix
Objective Realities	ix
Creative Capacities.....	x
Multiple-Application Workflows.....	x
Who Needs This Book?	x
How the Book Is Organized	xi
A Word to Windows Users.....	xi
This Book's Origins	xii
Image Credits	xii
Thank You!	xiii

Chapter One: What Is Sharpening? 1

Emphasizing Edges.....	2
Analog Roots.....	4
Digital Processing	5
Unsharp Mask Filter	5
Digital Noise Reduction	7
Sharpening and Pixels	9

Chapter Two: Why Do We Sharpen? 11

Sharpening and Displays.....	12
Display Technologies	12
Display Resolution	13
The Display as a Guide	13
Sharpening Issues	16
Image Sources	17
Detail, Film Grain, and Scanner Noise.....	18
Transparency Film Grain	18
Color Negative Film Grain	23
Detail, Noise, and Digital Capture	27
Sensor Types	27
Sensor Noise.....	33
Noise Reduction	36
Color Noise	36
Luminance Noise.....	37
Noise, Exposure, and ETTR.....	41

Noise, Detail, and Enlargement	45
JPEG and Noise	47
Lenses and Image Quality	52
Monochromatic Lens Aberrations	52
Chromatic Lens Aberrations	56
Lens Diffraction	58
Lens Flare	59
Photographic Aspects of Image Quality	59
Sharpening and the Image Source	64
Sharpening and Image Content	65
Low-Frequency and High-Frequency Images	67
How Sharp Is Sharp Enough?	71
Sharpening and the Output Process	74
Resolution and the Eye	76
Viewing Distance	78
Print Resolution	80
Halftone Output	80
Continuous-Tone Output	82
Error Diffusion Dither Output	82
Sharpening for Output	84
Introduction to Creative Sharpening	87
One Size Does Not Fit All	88

Chapter Three: Sharpening Strategies 89

Traditional Prepress Sharpening	90
One-Pass Sharpening Is Inflexible	90
One-Pass Sharpening Is Often Overdone	91
One-Pass Sharpening and Digital Capture	93
Traditional Sharpening Roots	96
A Multipass Sharpening Workflow	97
Phase One: Capture Sharpening	98
Phase Two: Creative Sharpening	102
Phase Three: Output Sharpening	104
Benefits of a Multipass Sharpening Workflow	108
Use-Neutral Capture-Sharpended Master Images	108
Fixed Output Sharpening	108
Creating a Multiapplication Sharpening Workflow	109
Sharpening in Photoshop	112
Sharpening in Camera Raw	113
Sharpening in Lightroom	114

Output Sharpening in Practice	116
Output Sharpening and the Display	118
From Theory to Practice	120

Chapter Four: Sharpening Tools 121

Sharpening Tools in Photoshop	122
Analyzing Sharpeners	122
The Sharpen Filter	123
The Sharpen More Filter	123
The Sharpen Edges Filter	124
The Unsharp Mask Filter	124
The Smart Sharpen Filter	134
The Overlay/High Pass Method	147
Edge and Surface Masks	152
Noise Reduction in Photoshop	159
The Reduce Noise Filter	160
The Surface Blur Filter	163
Noise Reduction with Despeckle	164
Comparing Noise Reduction Tools	164
Sharpening and Noise Reduction in Camera Raw and Lightroom	169
Comparing Camera Raw and Lightroom's Environments	170
The Camera Raw and Lightroom Sharpening Philosophy	172
Analyzing Camera Raw and Lightroom's Sharpening Tool	172
Using Camera Raw or Lightroom for Image Sharpening	177
Using Camera Raw or Lightroom for Noise Reduction	186
Combining Camera Raw, Lightroom, and Photoshop's Tools	189

Chapter Five: Industrial-Strength Sharpening Techniques 193

Layers, Blending, and Mask-Making Techniques in Photoshop	194
Sharpening on a Layer	195
Luminosity Blending	198
Sharpening Tonal Ranges	199
Sharpening Layer Controls	201
Splitting the Sharpening Haloes	202
Layer Masks	206
Capture Sharpening in Photoshop	214
Medium-Format Transparency	215
35mm Transparency	218
High-Speed 35mm Color Negative	220
Using Third-Party Raw Processing Software	223

Capture Sharpening in Camera Raw and Lightroom	227
Determining Edge Width.	228
Blending Image Frequencies with Smart Objects	233
Controlling File Size	241
Creative Sharpening in Photoshop	244
Choosing the Type of Mask	245
Sharpening Brush Techniques	247
Effects Layers	252
Creative Sharpening in Camera Raw and Lightroom	267
Lightroom to Photoshop to Lightroom for Creative Sharpening (and Other Things)	270
Output Sharpening in Photoshop	274
Halftone Sharpening	275
Inkjet Sharpening	276
Continuous-Tone Sharpening	277
Sharpening for Computer Display	278
Smart Sharpen for Display Sharpening	279
Output Sharpening in Camera Raw	281
Output Sharpening in Lightroom	283
Output Sharpening in the Print Module	283
Output Sharpening with Export	289
The Power of the Workflow	291
Chapter Six: Putting the Tools to Work	293
When Do We Sharpen?	294
Automation in Photoshop	294
Action Basics.	295
Batch Processing Using Photoshop	303
Saving Images with Image Processor	308
Efficient Capture Sharpening in Camera Raw	310
Creating Camera Raw Sharpening Presets	310
Controlling Camera Raw's Preferences	313
An Efficient Sharpening Workflow for Lightroom	316
Creating Sharpening Presets in the Develop Module	316
Creating and Using Templates in Print	321
Image Sizing and Output Sharpening in Export	328
Output Presets in Export	331
Post-Processing from Export	332
Sweating the Details	335
Index	337

PREFACE

Real World Sharpening

For almost 20 years, ever since it first shipped in 1990, Bruce and Jeff have labored in the vineyards of Photoshop, coaxing pixels to do our bidding. In that time, we've seen huge advances in the handling of image tone and color. When Bruce started digital imaging, getting the image to look the same on two different monitors was a major breakthrough, and matching the print to the screen appearance was the stuff of science fiction. Nowadays, it tends to be the rule rather than the exception.

We've made gargantuan strides in the handling of image tone and color, but when it comes to control of image *detail*, which is what this book is really about, we're still only a little more advanced than we were in 1990. It's just about impossible to make two different displays render image sharpness exactly the same way, and predicting print sharpness from the display is an exercise fraught with pitfalls and perils. In fact, we're back in the days of, "If you want to know what it will look like when you print it, print it and then look at it."

As a result, sharpening (and the equally important other side of the coin, noise reduction and smoothing) tends to be an *ad hoc* practice. We flail around until we get something that looks decent on the display and hope that appearance will somehow be transferred to the printed piece.

Photoshop CS4, Camera Raw, and Lightroom offer many powerful features for handling image detail. But how do you know what to aim for? This book contains a plethora of different sharpening tricks and techniques, but perhaps the most important contribution it strives to make is to provide you with an analytical framework that lets you think about sharpening in a new way. When we sharpen, we have to take at least three things into account:

- * The relationship between image detail and system noise that's imposed by the capture medium—we want to sharpen image details but we don't want to exaggerate film grain or digital noise.
- * The requirements of the image content—the wrong kind of sharpening can exaggerate texture we'd rather downplay, or even obscure detail instead of emphasizing it.
- * The needs of the print process—when we translate pixels into marks on paper (or canvas, or any other substrate on which we print) we inevitably introduce some softness for which we try to compensate by sharpening.

The huge problem is that these needs often contradict one another.

THE SHARPENING WORKFLOW

The solution is the sharpening workflow. By treating each demand separately, we can assure that all are addressed optimally. To some, this may smack of heresy: Doesn't everyone know that you can sharpen an image only once? We demonstrate that multipass sharpening is not only feasible but optimal. It's simply impossible to address all the conflicting needs of image source, image content, and output process in a single round of sharpening.

That said, multipass sharpening demands care and attention to detail. Blasting images with multiple hits of sharpening can create a hideously oversharpened mess, which is why the conventional wisdom dictates that you sharpen only once. However, the techniques described in this book allow you to sharpen images safely and optimally.

The sharpening workflow confers another benefit. By separating sharpening for output from the other sharpening processes, it creates use-neutral master images that you can easily repurpose for different output processes, at different sizes and resolutions.

Sharpening and the Display

One of the hardest sharpening lessons to learn is that what you see on your computer display can be highly misleading. But the screen is often all you have to rely on for your judgments before output.

Some display technologies render images much more sharply than others—the same image almost invariably looks sharper on an LCD display than it does on a CRT. Display resolution also has an impact. We’ve gone to some lengths to debunk the polite fiction that computer screens display images at 72 pixels per inch, and have even given instructions that will let you determine the real resolution of *your* display, which is the one that really matters.

The most important lesson of all, however, is that good sharpening for print can often look terrible—really, hideously, horribly bad—on screen. Learning the relationship between what you see on *your* display and what shows up on hard copy is a vital skill this book helps you acquire.

Objective Realities

Some parts of the sharpening equation are determinate. Human visual acuity—the ability to discern fine details—has limits that are rooted in the physiology of the eye. The same visual properties that we exploit to produce the illusion of continuous tone from dots of four colors of ink also have a direct bearing on sharpening.

Print sharpening is also a determinate process. Any given print process will always translate pixels into dots in the same way, regardless of image source or image content, so for any print process, there’s a right answer in terms of sharpening. (Of course, there are also many wrong ones.)

Creative Capacities

Sharpening is also a creative tool. We use sharpening to emphasize important detail (and sometimes we use blurring to suppress irrelevant, distracting detail), make a point, tell a story, invoke an emotion, or provide an illusion of three dimensions in our two-dimensional photographs. The sharpening workflow has a place for creativity, too. But it's important to know when to be creative and when to go by the numbers.

Multiple-Application Workflows

The first edition of this book dealt only with sharpening in Photoshop. For many reasons, a Photoshop-only workflow is no longer practical for most people using digital cameras. As a result of the changes in the industry (some of them helped by Bruce's philosophies) we wanted to describe and facilitate the approach of using the right tool for the job. Sometimes that won't be just Photoshop but may be Camera Raw or Lightroom or all three.

WHO NEEDS THIS BOOK?

If you work with images that are destined for hard copy, and you aren't totally confident about all your sharpening decisions, our hope is that you'll find this book beneficial. No matter whether you make your own prints, send them out to an online printing service, or deliver commercial work destined for offset press, the sharpening workflow can help you get the most out of your images.

This is not a book for Photoshop beginners, but neither is it a book only for Photoshop experts. Some of the techniques described herein use fairly esoteric Photoshop features with which you may or may not be familiar. Don't let that put you off. We've yet to encounter a piece of software that was smarter than its users, and Photoshop is no exception. Almost all the techniques in this book are nondestructive—they don't touch your original pixels—so you can't do any harm to your images by trying them.

How the Book Is Organized

We've tried to present all the information you need to build your own sharpening workflow.

The first two chapters look at the technical underpinnings of sharpening. Chapter 1, *What Is Sharpening?*, explores the fundamental nature of sharpening—what it does and how it works. Chapter 2, *Why Do We Sharpen?*, discusses the need for sharpening, and all the factors we need to address when we sharpen.

Chapter 3, *Sharpening Strategies*, provides an overview of the sharpening workflow and shows how it addresses each sharpening phase. Chapter 4, *Sharpening Tools*, describes the various tools that the sharpening workflow employs.

Chapter 5, *Industrial-Strength Sharpening Techniques*, shows how to use the tools described in Chapter 4 to satisfy the goals outlined in Chapter 3 by building a sharpening workflow from initial sharpening, through creative tweaking, all the way to final output.

Chapter 6, *Putting the Tools to Work*, shows how to create actions and do batch processing in Photoshop. It also emphasizes the importance of presets in Camera Raw and Lightroom. All of this is aimed at making your sharpening workflow efficient.

A WORD TO WINDOWS USERS

This book applies to both Windows and Mac OS. Since Bruce and Jeff have been using Macs for over 20 years, all the dialog boxes, menus, and palettes are illustrated using screen shots from the Mac OS version. Similarly, when discussing the many keyboard shortcuts in the program, we cite the Mac OS versions and try to include the Windows version. In every shortcut cited in this book, the Command key translates to the Ctrl key and the Option key translates to the Alt key. We apologize to all you Windows users for the small inconvenience, but because Photoshop is so close to being identical on both platforms, we picked the one we knew and ran with it.

THIS BOOK'S ORIGINS

Much of the material in this book is an outgrowth of work that Bruce Fraser originally did in the process of developing a commercial sharpening tool, PhotoKit Sharpener, from PixelGenius, LLC. Bruce was proud of PhotoKit Sharpener and was particularly proud to have been asked by Adobe to consult on the development of the advanced capture sharpening that went into Camera Raw 4.1 and Lightroom 1.1. While Bruce didn't get the chance to finish the consult, the Camera Raw team worked with Jeff to incorporate Bruce's capture-sharpening philosophies into Camera Raw and Lightroom.

PixelGenius is also proud that Adobe chose to work with us to adapt the PhotoKit Sharpener technology for Lightroom and Camera Raw. We appreciate that the Camera Raw and Lightroom teams have exhibited such respect to what Bruce thought was his most important contribution to the industry: the development of the concept of a sharpening workflow and its components; capture, creative, and output sharpening.

IMAGE CREDITS

Most of the images in the book were shot either by Bruce Fraser or Jeff Schewe. We are indebted to Martin Evening for the use of two of his images, which are credited in-place.

THANK YOU!

We couldn't have written this book without the excellent support of our publishing team, which includes Rebecca Gulick, Editor; Lisa Brazieal, Production Editor; Emily K. Wolman, Development and Copy Editor; Elle Yoko Suzuki, Proofreader; Owen Wolfson, Compositor; Valerie Haynes Perry, Indexer; Charlene Will, Cover Designer; and John Weber, Cover Illustrator.

Thanks to our pals and partners in PixelGenius, LLC—Martin Evening, Mac Holbert, Seth Resnick, and Andrew Rodney, for being the finest bunch of people with whom it has ever been our pleasure and privilege to work. An even bigger vote of thanks goes to the late Mike Skurski, our founding engineer, without whom we would never have been able to produce a successful software product. We all miss you. And thanks to the Pixel Mafia—you know who you are!

We also owe a debt of gratitude to the Knoll brothers (John and Thomas) for having the good sense to create Photoshop in the first place, as well as longtime engineer Mark Hamburg, with whom we've had the good fortune to work for many years. Our thanks also go out to the Lightroom team for their help and support.

And as always, Bruce would have wanted to thank his lovely wife, Angela, for being his best friend and partner, for supporting him in all his activities, and for making his life such a very happy one. Jeff thanks his wife, Rebecca, and daughter, Erica, for putting up with all the late hours (and missed family time) book writing seems to require.

This page intentionally left blank

CHAPTER ONE

What Is Sharpening?

And How Does It Work?

An old saw in photography goes, “If you want great prints, use a tripod!” While it’s usually delivered half-jokingly, the important grain of truth is that one of the ways our brains try to make sense of the world as seen through our eyes is by breaking down the scene into edges (objects) and non-edges (surfaces). If the edges in an image appear too sharp or not sharp enough, our brains tell us that there’s something wrong, and in the case of a photograph, the image appears unconvincing.

Sharpening is arguably one of the most important yet least understood aspects of digital image reproduction. Examples of badly sharpened images are easy to find—you probably need look no further than your daily newspaper or favorite magazine. Good sharpening, on the other hand, is invisible.

Sharpening can’t fix sloppy focus or a shaky camera. What it can and should do is to ensure that the sharpness of the original capture is carried through faithfully to the final output. Of course, sometimes we also use sharpening to improve reality—we may add some extra snap to the eyes and hair in a head shot, for example. But the primary purpose of sharpening is *not* to rescue overly soft images, but simply to counteract the inevitable softening that happens when we turn photons into pixels and pixels into marks on paper.

In Chapter 2, we’ll look at the various factors that give rise to the need for sharpening. But before we examine those, let’s look at how sharpening works, whether it’s done in a wet darkroom using analog tools or accomplished digitally either in the camera or in Adobe Photoshop, Adobe Camera Raw, or Adobe Lightroom.

Emphasizing Edges

Sharpening works by increasing the contrast around edges. Edges in images always involve darker tonal values adjacent to lighter ones. We can emphasize the edges by making the dark tonal values darker and the light tonal values lighter.

In the analog darkroom, this was accomplished using an *unsharp mask*. (See the sidebar “Why Is It Called ‘Unsharp Mask’ When It’s Used to Sharpen?” for details on just how this process worked.) In the digital domain, we sharpen by identifying the dark and light pixels that represent edges, and then we lower the value of the dark pixels to make them darker and raise the value of the light pixels to make them lighter.

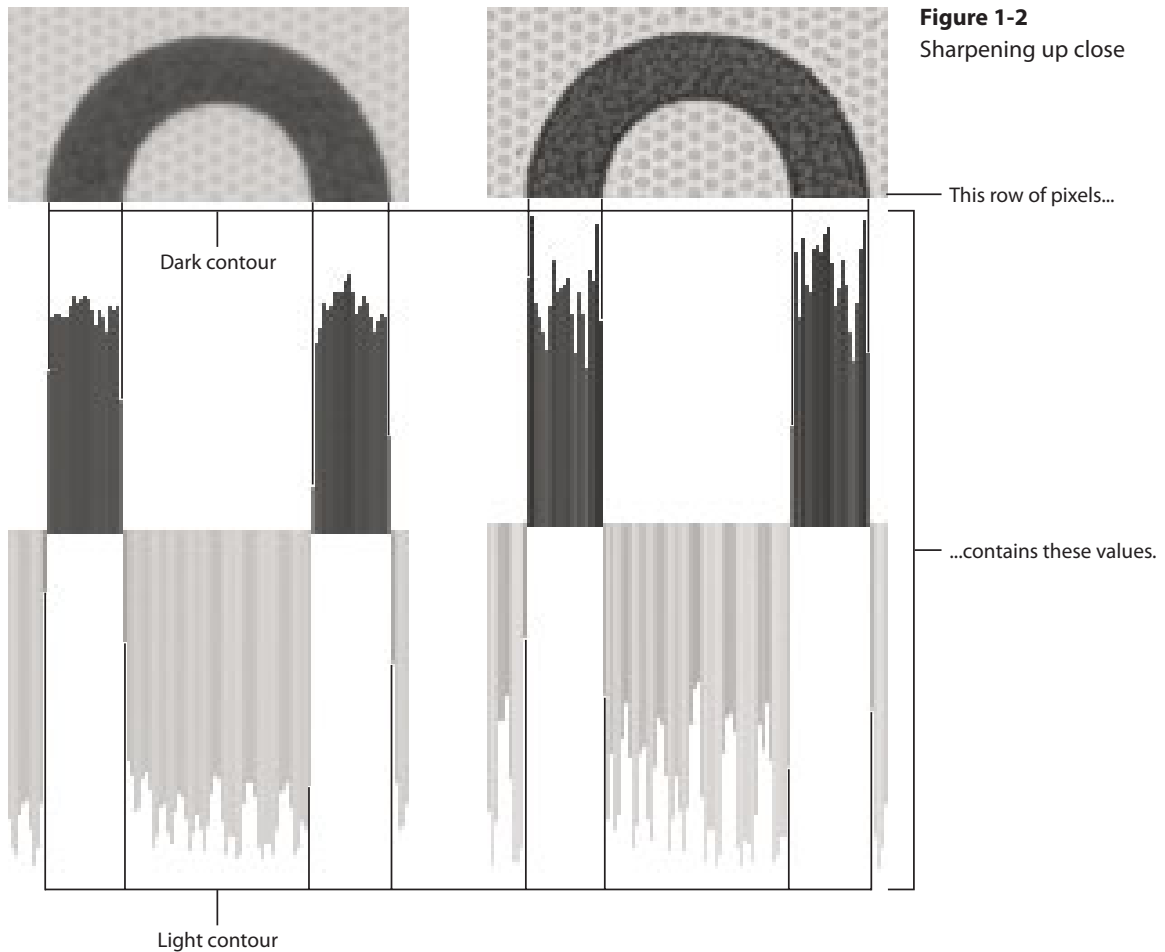
This adjustment creates a “halo” that makes the edges, and, hence, the entire image, seem sharper. The concept is simple, but as with many things in digital imaging, the devil is in the details that we discuss throughout this book!

Figure 1-1 shows the same image before and after sharpening. (The image also illustrates the pitfalls of driving in rural Scotland—a good metaphor for the myriad things that can go wrong when we use sharpening inappropriately.)

Figure 1-1 Before and after sharpening



The only difference between the two versions is the sharpening. Figure 1-2 shows a zoomed-in comparison with an accompanying graph of the values of a single row of pixels before and after sharpening.



Notice that the tonal range—the distance between the lightest and darkest tones—of the sharpened version is wider than that of the unsharpened version. Also notice that the biggest differences occur at the edge transitions of the capital O, while smaller differences emphasize the texture of the sign's background.

Sharpening is closely related to contrast, but simply increasing the contrast over an entire image just produces an over-contrasted image. Successful sharpening demands that we localize the contrast boost to those parts of the image that actually represent edges.

Why Is It Called “Unsharp Mask” When It’s Used to Sharpen?

Sharpening predates digital imaging by many decades. If you’ve often wondered why one of Photoshop’s main sharpening tools is named the Unsharp Mask filter when it’s supposed to make the image sharper, rest assured that you’re not alone.

The name originated from an analog wet-darkroom technique that increases the apparent sharpness of a photographic print using a duplicate of the negative to create a mask that increases contrast along the edges. The original and the duplicate negatives are placed on either side of a piece of glass—often just plain old window glass—and the entire sandwich is placed in the enlarger’s negative carrier.

When the enlarger is focused on the bottom negative, the top, out-of-focus copy creates a contrast mask that boosts the contrast along the edges in the image. Meanwhile, the out-of-focus dark contour burns the dark side of the edges and the out-of-focus light contour dodges the light side of the edges.

This technique is called “unsharp masking” because the mask—the top negative—is out of focus and, hence, isn’t sharp. In short, it’s an unsharp mask that has the effect of increasing the apparent sharpness in the print!

So the “unsharp” in “unsharp mask” refers to the mask, not the result, and it’s this analog technique that Photoshop’s Unsharp Mask filter replicates (though with much more control than its analog counterpart).

Analog Roots

As a photographic practice, sharpening has its roots in the analog world. However, the analog unsharp masking technique offered only two controls:

- The distance between the two negatives (the thickness of the glass) controlled the width of the sharpening halo.
- The exposure time controlled the strength of the contrast boost.

The tedious alignment, limited control, and uncertain results prevented unsharp masking from becoming a mainstream practice in analog photography. But it did see considerable use in sharpening analog color separations for offset printing. It was done by making a blurred duplicate of the continuous-tone separation and then printing it together with the sharp version

as a contact print, after which the separation was screened for halftone lines (lines of dots per inch).

Sharpening for continuous-tone photographic prints was something of a luxury, but when the continuous-tone original was turned into cyan, magenta, yellow, and black dots of ink, some of the original sharpness was lost. So sharpening became, and remains to this day, a standard operation in prepress. When the drum scanner replaced the stat camera (stationary camera used to shoot film from camera-ready artwork), digital sharpening became the norm.

DIGITAL PROCESSING

Digital sharpening tools offer far more precision, control, and options than the analog darkroom ever did. The inevitable downside is that with greater control comes greater responsibility, and the more options on offer, the more opportunities for mistakes, as we'll demonstrate throughout this book.


Unsharp Mask Filter

Photoshop's Unsharp Mask (USM) filter is a process that is similar to the analog version of sharpening. It creates an unseen (meaning you never see it) blurred version of the image in the background and uses the difference between the original version and the blurred version to determine edges that will be sharpened. Figure 1-3 shows a graphical representation of how Photoshop's USM filter works.

Figure 1-3 shows an original unsharpened image. The original layer was duplicated and the second layer was blurred three pixels using the Gaussian Blur filter. The blurred copy layer was set to Difference mode. To show the resulting difference, we've added a strong Levels adjustment layer to emphasize the results.

In the last image, the result of placing the layers in Difference mode was loaded as a selection (where white allows an effect and black resists) and used as a layer mask for a Curves adjustment to lighten the light edges and darken the dark edges. This is essentially the same function as the USM Radius and Amount adjustments, and is the basis of most digital sharpening algorithms.

Figure 1-3 Unsharp masking as simulated in Photoshop




Original image layer



Duplicated layer blurred three pixels



Blurred image over original image in Difference blend mode



Difference mode used as a mask to change contrast

The Threshold adjustment determines how much of a difference there must be between the blurred and unblurred copy before the sharpening process is applied. With a zero threshold, the sharpening would be applied to any difference. The higher the threshold, the less the sharpening will be applied for lower contrast edges. Figure 1-4 shows the result of increasing the threshold.

Increasing the threshold from zero to 50 means that only the higher contrast edges will be sharpened—the lower contrast edges will not.

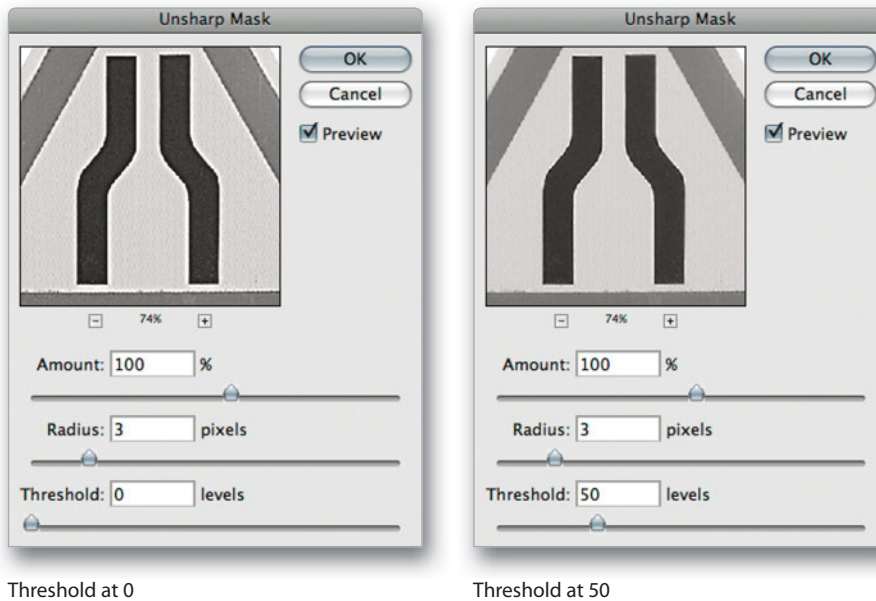


Figure 1-4 Threshold adjustment

Digital Noise Reduction

There are other digital sharpening algorithms above and beyond Photoshop's USM filter, and we cover them extensively in this book. However, there's one key difference between digital and analog sharpening with which we have to contend in varying degrees: the handling of noise. Analog sharpening rarely increased noise and often had the effect of reducing it, because the mask was made from a different piece of film than the original, so the distribution of the film grain—the main noise component in film images—was different. This had the desirable effect of making the noise in the mask cancel out the noise in the original. With digital sharpening, we aren't so lucky.

When it comes to digital sharpening, noise reduction is the flip side of the same coin. Sharpening tools such as Photoshop's USM filter evaluate one pixel at a time, determining how different it is from its neighbors. If the difference is great enough, the filter decides that this pixel represents an edge and adjusts its value accordingly.

The inherent problem is that the computer has no way of knowing whether the pixel differences represent real image information or noise, so unless we take steps to prevent it, most sharpening routines emphasize the noise as well as the edges. If you refer back to Figure 1-2, you'll see that the capital O in the sharpened version of the image is much noisier than in the unsharpened one.

However, if you refer even further back to Figure 1-1, it's unlikely that you'll see any objectionable noise in the printed reproduction at that size. The noise is confined to dark tones in small areas. Figure 1-5, on the other hand, shows what can happen if we fail to take noise into account.


Figure 1-5 Noise problems




Full image shows where the detail section came from.



This version was processed with no noise reduction and no sharpening.



This version was sharpened globally, emphasizing noise as well as detail.



This version had noise reduction applied before sharpening, and had sharpening applied through a mask.

This image is something of a “bad-case” scenario, shot at ISO 1600 on a fairly noisy digital SLR. It would have been a worst-case scenario if the shot had been unexposed, but it dramatizes the point that we must take noise into account when we sharpen if we don’t want our images to appear like sandpaper.

In many cases, we can prevent the sandpaper effect simply by protecting the noisy areas from sharpening. But in extreme cases like the one shown in Figure 1-5, some more proactive noise reduction may be needed before the application of sharpening.

Digital noise reduction is essentially the opposite of sharpening. Instead of increasing localized contrast, it reduces it, thereby rendering the noise less visible or even removing it entirely. Of course, the inevitable catch is that as it reduces the noise it also de-emphasizes the edges, making the image softer.

So one of the first key skills to master in sharpening is striking the right balance between emphasizing edges and reducing noise—walking the fine line between sharpness and sandpaper. Overly aggressive noise reduction creates an irreparably soft image, so noise reduction must be used with caution, and always should be done prior to sharpening (otherwise we sharpen the noise as well as the image). Most of the time, simply making sure that we don’t sharpen noise does the trick. But when images are captured on high-speed film or at high ISO settings on a digital camera, we may have to do a little more to reduce noise.

Sharpening and Pixels

All digital sharpening is carried out on pixels. Throughout this book, we’ll examine images at the pixel level to see what happens to the pixels through various stages of image reproduction. But, with the exception of images destined for on-screen display on the Web or in an information kiosk, the final product is not pixels but, rather, marks on paper or some other substrate.

The physical relationship between the image pixels on a display and the marks on paper is rarely obvious. In conventional halftone printing—the kind we use on printing presses—it typically takes between two and four pixels to form a single halftone dot in grayscale printing. When we print in color, this relationship is even less obvious, since the halftone dots are spread across four different color plates.

Inkjet printers use a different type of screening, but again, the relationship between printer dots is complicated. Even in the relatively simple scenario of printing to a true continuous-tone printer, where the pixels are reproduced one-to-one, the pixels on the print are typically much smaller than the same pixels on screen.

We can't control how the pixels get turned into printer dots, yet we can control the degree of contrast boost and the size of the pixels on output. In so doing, we control, in turn, the width of the sharpening halo. The goal is to avoid haloes that our eyes pick out as discrete features, while making them strong enough and just wide enough to produce a sharpening effect.

Since we can control only the final size of the pixels, we must do our final sharpening at final output size and resolution to get the sharpening halo right. If the image is resized after sharpening, the haloes either get bigger and visually obvious, or they become too small and we lose the sharpening effect.

For some people, that means that *all* sharpening must be deferred until the image is at output size; otherwise, multiple rounds of sharpening will destroy the image. That's one of the notions this book seeks to challenge: It's certainly true that multiple rounds of sharpening *can* destroy the image, and that the final, critical round of sharpening must be tailored to the output process and applied to the final image pixels. But throughout this book, we argue that optimal sharpening simply cannot be accomplished by a single pass of sharpening at the end of the image preparation process. This is because the softening that occurs when pixels are turned into print dots is only one of several, often competing factors that must be addressed.

In the next chapter, we'll look at all of the different factors we must take into account when we sharpen.

CHAPTER TWO

Why Do We Sharpen?

And What Must We Take Into Account When We Do So?

Whenever we turn photons into pixels, we lose some degree of sharpness. No matter how high-resolution our capture devices or how expensive our lenses, they sample a fixed grid of pixels and are thus limited. So cameras (and scanners) turn the continuous gradations of tone and color that exist in the real world into discrete pixels of the digital world. When the pixels are small enough, they provide the convincing *illusion* of continuous tone, but it *is* just an illusion. Image sharpening is one of the things we simply must do to make the illusion convincing.

The reason that sharpening is such a complex topic (and the reason for this book) is that successful sharpening requires the consideration of several potentially competing factors. These factors, which we'll discuss in detail in this chapter, fall into three basic categories:

- The image source
- The image content
- The image use

For decades, the standard operating procedure has been to punt on the issue of how and when to sharpen, and just try to handle all sharpening in one single pass, usually at the end of image processing, just before output. With film scans, this worked after a fashion, though we suspect the results were rarely optimal. With most digital captures, the results are even less likely to be optimal simply due to how most digital cameras work.

Does My Image Need Sharpening?

Yes. Even if your image appears reasonably sharp when viewed at Actual Pixels (the Adobe Photoshop View menu command that displays the image at 100% zoom) on your computer display, the image will need sharpening depending on the final use. Of course, if your image appears in soft focus in Actual Pixels view, it definitely needs sharpening no matter what the final use.

One of the harshest sharpening lessons we've learned is that it's extremely difficult to make judgments about output sharpening by looking at pixels on a computer display. If your images are destined for a specific type of display—for example, a kiosk—ideally, you should look at the images on that display before making any final judgments. Barring that, use the display as a guide, as explained in the “The Display as a Guide” section, below.

SHARPENING AND DISPLAYS

 **NOTE:** You'll notice that we refer to the screen that connects to your computer that allows you to see images as a “display.” A computer “monitor” is something that monitors your computer, not a display. We realize this is a small distinction, but we have good friends in the computer industry who would complain if we referred to the device on which we view images as a “monitor.”

Unless a TV or computer display is your final output, making images look sharp on the display almost invariably results in an under-sharpened printed output. But even displays themselves vary hugely in the way they reproduce sharpness.

Display Technologies

LCDs (liquid crystal displays) are much sharper than CRT (cathode-ray tube) displays when set to their native resolution. An image that looks sharp on an LCD device may still look soft on a CRT, and an image that looks sharp on a CRT display may look slightly crunchy on an LCD.

If your images are being prepared for the Web, realize that you have very little control over what the other billion or so Web users will see. Nowadays,

many more users have LCDs than CRT displays. In fact, CRTs are no longer being made. All you can do is aim for the middle of a very large barn door, but remember that an LCD is likely to be more representative of the general Web population than is a CRT.

Display Resolution

The resolution of the display also has a major impact on how we see image sharpness. An image that looks pleasingly sharp on a display running at 1600 x 1200 shows obvious, less-pleasing sharpening haloes on the same display running at 1024 x 768. The pixels are displayed at a larger size when the display is run at 1024 x 768, so the sharpening haloes become bigger and, thus, more obvious.

Again, if you're sharpening images for the Web, you need to aim for the "average" display resolution, which is probably closer to 1024 x 768 than 1600 x 1200.

The Display as a Guide

Sharpening for use on a display is an uncertain endeavor unless you know the exact display for which you're sharpening. Yet a much bigger issue is how to use the display as a *guide* for print sharpening. The easy answer is, unfortunately, that there are no easy answers!

Luckily, the situation isn't entirely hopeless—it's just challenging. The fundamental concern that should always be foremost in your mind is the actual size of the pixels on output. But to understand the relationship between what you see on the screen and what you get on output, some additional information is useful.

First, determine the display's true resolution. Many of us tend to leave unquestioned the polite fiction that computer displays have a resolution of 72 pixels per inch (ppi). This is often not the case—in fact, it's quite unlikely that your display's resolution is exactly 72 ppi, though it may be close.

Display vendors usually specify the size of the image area as a diagonal measurement, which makes for a big number but isn't the actual resolution. Determining the display's true resolution is a simple exercise that requires no equipment more complicated than a tape measure:

1. Using a tape measure, measure the *width*, in inches, of the image area on your display.
2. Divide the number of horizontal pixels that your display can show by the number of inches from Step 1.

The resulting number is the true resolution of your display in pixels per inch.

For example, the horizontal image area of both of Jeff's main LCDs (a 30" NEC 3090WXi and an Apple 30" Cinema) is 25.25 inches. The native resolution for both is 2560 x 1600 pixels, so the actual resolution is 101.4 ppi. A third LCD has a resolution of 1920 x 1200 pixels and is running at 87.8 ppi.

The image area on CRTs is more variable than on LCDs, because the geometry controls let you adjust the picture size.

Table 2-1 shows the approximate resolutions for typical display sizes and pixel dimensions.

Table 2-1 True display resolution

Display size	Pixel dimensions	Resolution (ppi)
30" LCD	2560 x 1600	101
21" LCD	1600 x 1200	95
	1280 x 1024	76
	1024 x 768	61
	800 x 600	47
21/22" CRT	1920 x 1440	122
	1600 x 1200	102
	1280 x 1024	82
	1152 x 870	74
	1024 x 768	65
	800 x 600	51
17" laptop	1920 x 1200	133
15" laptop	1440 x 900	100

Knowing your display's true resolution is the key to understanding the relationship between the pixels you see on your display and the final printed results.

If your display is around 100 ppi and you're printing at around 240 ppi, viewing the image at 50% will give you a truer picture of the final sharpening. Similarly, if your display is around 75 ppi and you're printing at 300 ppi, the 25% view will give you a closer idea of final sharpness than any of the higher zoom percentages. It still won't be perfect—Adobe Photoshop's downsampling algorithms, regardless of the version, have a different effect on sharpness than do the mechanisms by which printers or platesetters turn image pixels into dots—but it'll be a lot closer than the Actual Pixels (100%) view. See the sidebar “How Does Photoshop Display Images?” for information about changes in Photoshop CS4's display engine.

Thus the problem remains: You simply cannot reliably judge image sharpening intended for printed output on a computer display. Even if you reduce the screen zoom, you are still evaluating your higher-resolution image on a low-resolution device. Neither the image sharpening nor image noise nor grain can be displayed accurately at a reduced zoom. The effective resolution of the display will be a third to a quarter of the resolution contained in the image.

How Does Photoshop Display Images?

Photoshop CS4 takes the open image (or images) and scales it down (downsamples) from its actual resolution to a size and resolution for on-screen display. In previous versions of Photoshop, users had to avoid odd zoom percentages because Photoshop scaled images for display using a Nearest Neighbor interpolation for speed. The 50%, 25%, and conceivably the 12.5% views preserved sharpness reasonably well, but the odd zoom percentages (66.7%, 33.3%, and intermediate zoom percentages) did not, because Photoshop applied fairly strong anti-aliasing (partial filling of edge pixels to create smooth line edges) to those zoom levels. Photoshop CS4's ability to use your computer's video card's GPU (graphics processing unit) for display enhancements is a major change in how Photoshop can display images.

With the Enable OpenGL Drawing option checked, Photoshop uses high-precision modeling and smoothing by using a bilinear interpolation when scaling images for on-screen display. However, while this is a great improvement in display quality, it further complicates an already complicated situation, because not everybody has suitable video cards, or users may need to turn off the Enable OpenGL Drawing for compatibility reasons. And while it certainly can make a wide range of display percentages look good, the GPU doesn't magically add any resolution to your display. So your computer's display is still a low-resolution viewing device (even if it looks a lot better in Photoshop CS4).

The only really reliable method of actually seeing your image pixels is to evaluate the image at a screen zoom of 100%, where one image pixel will be displayed on one display pixel. Only at this 1:1 ratio can you see what image detail is contained in the image. The problem with this is that at 1:1, the image will be two to four times larger than the printed dimensions. At that size, the relationship of image detail to printed detail is inaccurate. Thus there are no good solutions short of making a print and evaluating it.

The key point to all this is that when we look at sharpening on a display, we're viewing it "through a glass, darkly" (though the Apostle Paul doubtless had something else in mind when he coined the phrase). At some fairly far-off future date, we may benefit from some technology that compensates for the large variation in the ways different displays show the same pixels' sharpness, just as color management currently lets us compensate for the way different displays show the same pixels' colors.

But that day won't come soon. Until then, you'll have to learn to make that compensation in your head. This is not trivial, but is it not impossible. First, you need to understand the size relationship between your pixels on screen and your pixels on output. Then it's a matter of learning the behavior of your display, and making constant comparisons between what it shows you and what ends up in print, so that you gain experience as to what will work. Throughout the remainder of this book, we'll provide pointers as to the kinds of things to look for.

SHARPENING ISSUES

As photographers and digital imaging artists, we are confronted by a variety of sharpening issues. Most prominent is the inevitable softening that occurs whenever we digitize, or turn a continuous stream of photons into discrete square pixels. No matter how good our lenses, no matter how high-resolution our digital cameras or scanners, the digitization process always loses some apparent sharpness.

In addition, there are the normal photographic aspects that impact sharpness, such as camera or subject blur, lens diffraction, optical defects, and issues of film grain or sensor noise.

Further softening occurs on output, when we translate image pixels to printer dots. Each output process has its own requirements for sharpening—we must apply very different sharpening to the same pixels printed at 300 ppi on a continuous-tone printer than at 300 ppi on a 150-line screen halftone printer.

Last but not least are the issues presented by the image content itself. A forest of trees needs different sharpening treatment than a model's headshot. Sharpening that emphasizes the fine detail in the trees turns the model's skin into a moonscape, while sharpening that emphasizes the model's eyes and lips may obscure the detail in the forest. Somehow, we have to reconcile the various and often contradictory needs of the image source, the image content, and the output process.

We'll spend the rest of this chapter looking at these disparate demands in detail, because until they're understood and addressed, sharpening won't do justice to the image—and may actually do more harm than good.

IMAGE SOURCES

The image source imposes its own signature and its own limitations on the image. The signature is the relationship between the noise that we want to de-emphasize and the detail that we want to emphasize. The limitations are the degree to which the image can be enlarged and the smallest details that can be captured.

In part, enlargement is limited by the gross pixel count, especially with digital captures. But the noise signature also plays a role—we typically scan film to obtain a significantly higher pixel count than digital captures offer, yet film scans cannot withstand as much enlargement as a digital capture of equivalent pixel dimensions. This is because digital captures generally have less noise than scanned film: The interaction between the film grain and the scanner's pixel grid exaggerates the grain to produce a noise signature that's stronger than noise from digital cameras.



NOTE: Much more heat than light has been generated in discussions of the equivalent information content of digital and film captures. Without adding too much fuel to the fire, our experience tells us that 35mm film scanned at 6300 ppi (creating a file approximately 8900 x 5700 pixels) has about the same potential for enlargement as a digital capture in the 6 to 8 megapixel range (approximately 3000-3500 x 2000-2350 pixels). This is a very approximate rule of thumb.

DETAIL, FILM GRAIN, AND SCANNER NOISE

A primary sharpening challenge lies in emphasizing detail without emphasizing noise. We can often accomplish this by simply protecting the noisy areas from sharpening. In more extreme cases, we may need to apply some noise reduction prior to sharpening. (We'll discuss specific techniques for doing both in Chapter 4, "Sharpening Tools.")

Many of the techniques you'll need to deploy depend on the final use of the image. Downsampling can hide a great deal of noise, making the image easier to sharpen, so for one-off uses at sizes that are small enough to allow downsampling by 50% or more, noise may be less of a concern. But as we approach the practical size limitations of the capture, the noise becomes ever more obvious and may need special handling.

Scanned transparencies, scanned negatives, and digital captures all have their own noise characteristics. They share the property that noise increases with higher ISO (International Standards Organization) ratings, but they differ in important ways too. First, let's look at scanned film.

Transparency Film Grain

The grain in scanned film becomes visually obvious at scanning resolutions far below that required to resolve individual grains. (Resolving individual grains would require scanning at ridiculously high resolutions of around 11,000 ppi, providing no more usable image information than a lower-resolution scan.)

Grain is most obvious in small formats like 35mm, and much less obvious in medium or large-format transparencies. However, much of the image noise comes not from the grain itself, but from the interaction of the grain, the scanner's sampling grid, and any digital noise introduced by the scanner.

The grain in color transparencies tends to be monochromatic, but the scanner noise is typically strongest in the blue channel, since the blue filter is the least efficient in terms of allowing light to pass through. As a result, transparency grain often shows up at its strongest in clear blue skies, though close examination of the pixels shows that the grain is present in the other channels, albeit to a lesser extent.

Figure 2-1 shows three views of typical scanned 35mm transparency noise:

- The first view is an attempt to show the actual image pixels, printed at 72 ppi. Note, however, that the grain viewed on screen actually looks worse than the grain in print due to the softening effect of the halftone process.
- The second view is a “contact print” that reproduces the actual scanned pixels at 300 ppi, the optimal resolution for the 150-line screen with which this book is printed. This view represents the optimal enlargement one can reasonably attempt from this capture. (On many images, we can drop the resolution to 225 ppi, and hence make a bigger print, but fine diagonal features will usually suffer.)
- The third view shows the entire image, still unsharpened, but downsampled to 300 ppi. The image appears reasonably sharp and more or less grain-free, due to the massive downsampling—the image is now only 16% of the original captured size.

Figure 2-2 shows noise reduction and sharpening for the image at the contact print resolution. The sharpened version was created by applying noise reduction and preliminary sharpening, and selectively using layer masks to isolate the edges on the full-resolution 6300-ppi capture. The image was resized (but not resampled) to 300-ppi, and then sharpening was applied, which was tuned to the halftone printing process.

For the downsampled image, the preliminary sharpening and noise reduction for the optimally sharpened version were performed on the full-resolution 6300-ppi capture. Then the image was downsampled to 300 ppi using Photoshop’s Bicubic Sharper algorithm and sharpened for the halftone process.

The differences between Figures 2-1 and 2-2 are much more subtle on the downsampled version than they are in the contact print, reinforcing the point that final output size makes a big difference. Also, note how crunchy and oversharpened the zoomed-in view appears. If one were sharpening by eye, it’s unlikely that one would be willing to take the sharpening and noise reduction this far.

Figure 2-1 Noise in scanned transparency

This zoomed-in view shows the unsharpened pixels from a 35mm transparency scanned at 6300 ppi, printed at 72 ppi.



This view shows a larger section of the unsharpened image printed at 300 ppi to produce the maximum size the capture can support optimally. (The full-sized image would be approximately 20" x 30".)



This view shows the entire image downsampled to 300 ppi, with the zoomed area above outlined in black. The image is still unsharpened, but the noise is largely rendered invisible due to the downsampling.





Figure 2-2 Scanned transparency with noise reduction and sharpening

This zoomed-in view shows the pixel crop from the scan printed at 72 ppi. This has undergone noise reduction, selective sharpening, and final output sharpening. At this zoomed-in view, the image appears over sharpened.



This view shows a larger section of the reduced noise and sharpened image printed at 300 ppi. This is the output size for which the noise reduction and sharpening was intended.



This view shows the entire image downsampled to 300 ppi with noise reduction and sharpening applied. Note that the output sharpening was applied after downsampling using Photoshop's Bicubic Sharper interpolation algorithm.

Figure 2-3 shows the individual color channels of the unsharpened scan in the same Actual Pixels view as the first illustration in Figure 2-1, and how dramatically noisier the blue channel is than the red or green. This is more a function of scanner noise than film grain. The blue tricolor separation filters in most scanners are terribly inefficient and absorb a lot of light, so less light is gathered in the blue channel of the film scan. The purpose of the figure is to convey an impression of typical grain in a transparency scan, and to show its diminishing influence as we reduce the image size.

Figure 2-3 Noise in scanned transparencies by channel

This zoomed-in view shows the unsharpened pixels from the red channel.



This zoomed-in view shows the unsharpened pixels from the green channel.



This zoomed-in view shows the unsharpened pixels from the blue channel. Notice how much more obvious the noise in the blue channel is compared to the red or green.

