



INTRODUCTION TO GRAPHICS

Sharpening Workflow

Information
Sheet No. PS958

Sharpening is a critical step in digital imaging, yet few people are satisfied with the results. Research conducted for a commercial sharpening product, produces a new approach to sharpening that's based on three imaging phases: capture, creative, and output.

In most situations users of Photoshop are encouraged to use a two-pass approach to sharpening.

Sharpening is one of the most-overlooked yet most-essential components of the digital-image-reproduction workflow. People often ask when to sharpen, and what settings to use: The truth is that neither of these questions has a single, let alone a simple, answer.

Why Sharpen?

You need to sharpen for several reasons, each of which imposes its own demands. But more often than not those demands contradict each other.

Whenever you turn photons into pixels, you lose some sharpness, because no matter how high the resolution of our capture devices, they sample a fixed grid of pixels, turning the continuous gradations of tone and color that exist in the real world into discrete pixels.

Sharpening can be a creative tool. Sometimes you may want to make the image sharper than it really was, to tell a story, make a point, or emphasize an area of interest.

When you turn pixels into halftone dots or stochastic dithers of inks, or even into continuous-tone dots on dye-sublimation printers or film recorders, you will lose sharpness again. So our output devices also introduce some softness.

There is a need to apply sharpening to counteract the softness introduced by both capture and output, and to satisfy yourr creative goals.

The Sharpening Problem

The problem inherent in trying to achieve all sharpening goals in a single pass is that we need to sharpen images for at least three basic reasons, and they often present conflicting demands. Why is this?

The capture process, whether scanned film or direct digital capture, imposes its own noise "signature" on the

image. You want to sharpen the detail without exaggerating the noise, so successful sharpening needs to take into account the relationship between image detail and the noise signature of the image source, whether film grain or a digital camera filter mosaic. Incorrect sharpening makes the grain or noise more rather than less obvious, and sometimes even obscures the detail instead of enhancing it.

Good sharpening is content-sensitive. Apparent sharpness depends on the contrast along what you see as edges. A close subject with soft detail such as a headshot has wider edges than a high-frequency image such as a forest full of trees. If we use the same sharpening for both, the results will be less than optimal on at least one of the images. Incorrect sharpening either obscures small details, or over sharpens textured areas such as skin tones.

The output process introduces softness too, and you need to sharpen to counteract it. But the sharpening required for a traditional ink-on-paper halftone may be fairly different from that required for an inkjet using error diffusion rather than a conventional halftone, and different again for a true continuous-tone printer used for digital prints. Incorrect sharpening either produces insufficiently sharp results, or makes obvious, objectionable sharpening haloes along high-contrast edges.

The Case for Three

You will most certainly never be happy with the results when you try to address all these problems in a single round of sharpening. Over the years, Most people have flirted with a two-pass approach to sharpening, performing a round of gentle sharpening close to capture time, and a second round tailored for the specific output process, but it was going as much by instinct as anything else.

Since you sharpen for three very different reasons, why not split the sharpening into three separate stages: one close to capture, one in the creative phase, and one round of sharpening tailored specifically for the final output? The obvious reason not to do so is that you'll end up with a hideously over sharpened mess. But based on extensive testing and research, it's not only possible, but also optimal to break up sharpening into these three stages. Obviously, caution is required, but before you

look at how to make such a sharpening workflow, let's look at the advantages.

First, applying a modest amount of sharpening close to capture time makes it much easier to make good judgments about fine-tuning contrast. Sharpening is essentially a localized contrast adjustment — increasing the contrast along edges — and global Unsharp Masking can often change the tonal balance of an image in unintended ways. Nudging the image towards reasonable sharpness early on helps the editing process, and provides a solid floor to stand on when it's time to make creative sharpening decisions.

Second, uncoupling capture and creative sharpening from sharpening for output lets you keep a single master file to which capture and creative sharpening have already been applied, and from which can be used to create different iterations for different print processes. Once the particular version is at final output size, you can apply an output sharpen that's specifically tailored for that output process.

The Sharpening Workflow Solution

The Big Three factors that sharpening has to address, bears a stronger resemblance to a color management problem. At the risk of drawing a parallel from one incomprehensible subject to another, you could make a strong analogy between the sharpening workflow and a color management workflow.

The conversion from a capture color space to a Photoshop working space for editing is analogous to applying a source-specific capture sharpen. The tone and color edits you make in the working space are analogous to applying creative sharpening. The color conversion you make to an output profile for output on a specific device is analogous to output sharpening. The master image contains "device-independent" sharpening that holds up well through resizing operations, and the final "device-specific" sharpen is applied to the output file at output resolution.

Translating this workflow concept to sharpening is actually rather simple. Instead of trying to handle all the different issues that affect sharpening in a single edit, the sharpening workflow splits sharpening into three stages:

1. Capture Sharpening is applied early in the image-editing process, and just aims to restore any sharpness that was lost in the capture process.

2. Creative Sharpening is usually applied locally to accentuate specific features in an image—for example, you often want give eyes a little extra sharpness in head shots.

3. Output Sharpening is applied to files that have already had capture and creative sharpening applied, after they've been sized to final output resolution, and

is tailored to a specific type of output process.

Dividing sharpening into these three phases lets you handle the competing needs separately. You need to compensate for the image source and image type using capture sharpening, without having to worry about output issues, and apply creative sharpening directly to the image at hand without having to worry about image sources (since they've already been addressed by the capture sharpener). So output sharpening can then concentrate solely on the needs of the output process since we've handled the other problems in the two previous phases.

Implementation and Techniques

Creating a sharpening workflow in Photoshop, and making that workflow work, requires patience, skill, and a willingness to toss out any cherished dogmas. In the early stages at least, it also demands a willingness to print images that look downright scary on screen — more on that later.

Capture Sharpening. Capture sharpening must be done carefully: apply Unsharp Mask to a pixel layer in Luminosity mode, using an edge mask either as a layer mask, or as the selection for the sharpened pixels. That way, the sharpening is non-destructive and the luminosity blending prevents color fringing. You can modify the global sharpening strength by varying the opacity of the sharpening layer, and modify the local sharpening by changing the contrast of the mask.

Building edge masks is relatively simple. Start out with a channel that has good contrast along the edges you want to emphasize — sometimes you can simply duplicate one of the color channels, but Photoshop's Calculations feature offers a formidable array of options for obtaining a channel with the necessary contrast.

Once you have the channel, run the Find Edges filter (lurking in the Filter>Stylize menu). It does a decent job of finding the edges, but you need to invert the result for the mask — you want the edges to be white, so that they'll get sharpened, and the non-edges to be black, so that they won't. Add a couple of steps to soften the transition between sharpened and non-sharpened areas, using Gaussian Blur to soften the contours and Levels or Curves to adjust the contrast of the mask. Figure 1 shows the unsharpened image and the edge mask.



FIGURE 1: THE UNSHARPENED IMAGE (LEFT) AND EDGE MASK (RIGHT).

Using the edge mask as a selection, you can copy the selected pixels to a new layer set to Luminosity blending, or you can copy the entire image to a new layer, also set to luminosity blending, then apply the edge mask as a layer mask. The latter approach gives you more control after the sharpening has been applied, but at the cost of a larger file and hence slower processing.

At this stage, it's important to treat the highlights and shadows very gently, otherwise they'll get blown out and plugged up, respectively, in the subsequent sharpening rounds. The Blend If: sliders in the Layer Options dialog box are invaluable for constraining the capture sharpening to the midtones -- you can access these by choosing Layer Style>Blending Options from the Layer menu, or by simply double-clicking on the layer in the Layers palette. Figure 2 shows the layer blending options set to protect highlights and shadows, focusing the sharpening on the midtones.

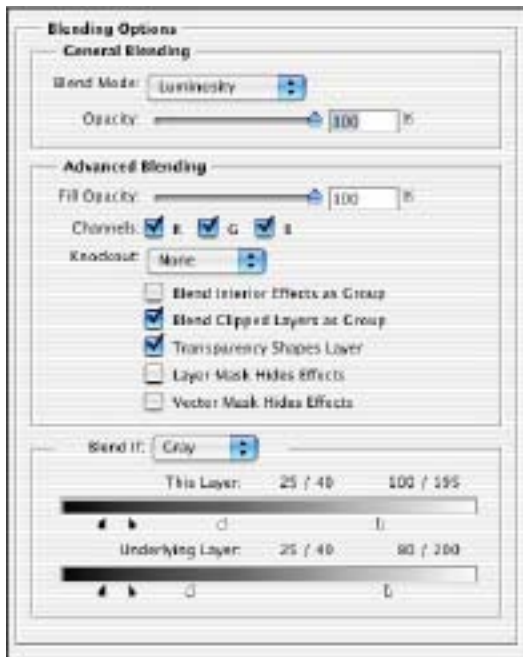


FIGURE 2: LAYER BLENDING OPTIONS.

The result is a clean, gentle sharpen with no obvious haloes. Figure 3 shows a detail of the image unsharpened, and with the capture sharpening applied.



FIGURE 3: DETAIL UNSHARPENED (LEFT) AND SHARPENED (RIGHT)

Creative Sharpening. The only real guideline that can be give here is to use common sense. One useful way to apply creative sharpening is to make a new layer and merge the visible layers into it (Control-Alt-Shift-E), set the blend mode to luminosity, apply a somewhat stronger Unsharp Mask than is ultimately desired, with a radius setting that picks out the details you want to emphasize, then add a layer mask set to Hide All.

Use a low-opacity brush set to white, painting on the layer mask, to paint the effect into the areas that need it. When a layer mask is targeted for editing, the foreground and background colors are automatically set to black and white. You can easily toggle between painting with white to paint the effect in and painting with black to paint the effect out by pressing the 'x' key.

Figure 4 shows the image detail after applying a creative sharpening layer created as follows. Create a new layer, then Option/Alt Merge Visible to copy all the pixels to the new layer. Set the Layer Blending Options to Luminosity at 66% opacity (so that I can increase or decrease the effect after the fact), run Unsharp Mask with Amount=500, Radius = 0.3, Threshold=0, add a layer mask set to Hide All, then paint the sharpening in with a brush set to 50% opacity.



FIGURE 4: DETAIL AFTER CREATIVE SHARPENING, WITH CREATIVE SHARPENING MASK.

Output Sharpening. In the capture and creative sharpening phases, you want to avoid obvious sharpening haloes. With output sharpening, though, that's not the case. The rule of thumb is to aim for a sharpening halo that's somewhere between 1/50th and 1/100th of an inch in width -- the thinking being that at normal viewing distances it's too small to see as an actual halo. It follows that output sharpening must be done at final output resolution.

To accomplish this, you need to know the output size and resolution, and the way in which the output device converts the pixels into marks on paper. For example, if you're printing to an inkjet printer at 300 ppi, you want to create a sharpening halo that's about 3 pixels wide. If you're printing to a halftone printer such as a press, using a 266-ppi file and a 133-line screen, you need a halo that's at least 4 pixels wide, because each halftone dot is made up of four pixels, and you may well be rewarded if you make the halo 5 or 6 pixels wide. Output sharpening doesn't require any complex masking -- all the image-specific local concerns have already been addressed in the capture and creative sharpens -- so the output sharpening can be applied globally.

The result looks pretty ugly on screen when viewed at an Actual Pixels zoom level. Good advice is to make the file look "crunchy" on screen when sharpening for offset printing: Research suggests that you can make better judgments about output sharpening by viewing at 50% or even 25% (avoid the "odd" percentages like 66.6% because Photoshop applies heavy antialiasing to those views), but the only truly reliable way to make the judgment is to view the printed result. If you follow the above-mentioned formula, keeping the haloes to between 1/50th and 1/100th of an inch on the final output, the results will look good.

For output sharpening, use a duplicate of the image set to Overlay, then run the High Pass filter. Set the Layer Blending options to protect the extreme highlights and shadows -- there doesn't seem to be much point in sharpening pixels brighter than level 245 or so, or darker than level 15 or so.

Figure 5 shows the same image detail sharpened for inkjet output at 300 ppi. On screen, it's quite crunchy, but the printed result shows no obvious haloes, just a sharp print. You can use this same approach for all kind of output, simply varying the High Pass radius to suit the output resolution.



FIGURE 5: DETAIL AFTER OUTPUT SHARPENING FOR INKJET AT 300 PPI.

The monitor is also an output device. You know that when you drastically downsample an image to screen resolution for display on the Web, you need to resharpen it. Figure 6 shows the original 3072 x 2048-pixel image downsampled to 400 x 600 pixels, including the capture and creative sharpening layers. It's rather soft, though not objectionably so.



FIGURE 6: DOWNSAMPLED IMAGE INCLUDING CAPTURE AND CREATIVE SHARPENING.

But the sharpening layers make the image very responsive to a simple output sharpen for Web resolution, again using an Overlay layer filtered with the High Pass filter, this time at a 1-pixel radius. Figure 7 on the next page shows the downsampled image with output sharpening applied for Web resolution.



FIGURE 7: DOWNSAMPLED IMAGE SHARPENED FOR WEB RESOLUTION.

Looking Sharp

These ideas may not solve all the sharpening problems in the digital universe, but breaking up the sharpening process into individually addressable chunks makes total sense from the standpoints of both quality and workflow. In particular, treating output sharpening as a device-independent process, separate from the needs of both the image and the capture mechanism, allows you to sharpen optimally for each output process. This tutorial at least represents the first baby steps towards that goal.

On (too) many occasions, you will be asked to prepare files for output without knowing the output process. Ultimately the solutions to this problem may be political rather than technological, but experience suggests that images with careful capture and creative sharpening will hold up much better to unknown downstream sharpening than either an unsharpened image, or one that has been sharpened by guesswork for unknown output. In the long term, it may make sense to apply output sharpening in the RIP, just as many RIPs now perform a final color conversion.

A workflow approach to sharpening also lends itself to automation. Capture sharpeners can be designed to accommodate the needs of specific capture

devices -- the inclusion of metadata in the vast majority of digital camera files will be an important aid to this -- while different image types may be addressable in the short term by a series of presets, and in the long term by performing some kind of frequency analysis to determine the dominant tendency of important edge details. Output sharpeners, like output profiles, are then tailored purely to the needs of the specific output device, since all the image-specific concerns have already been dealt with.

At the very least, a workflow approach to sharpening gives a much better conceptual framework in which to cast sharpening issues than simply trying to find a one-size-fits-all approach (which is simply impossible). For you the real payoff is that sharpening becomes a creative tool: The capture sharpen addresses the defects of the capture; the output sharpen addresses the defects of the output. In the middle, there's lots of room to play.