

# THE MAKEREADY ARCHIVE

Column **69** of 77

## Life on the Edge

*Topics:* Unsharp masking, hiraloom sharpening, Gaussian blur filter, flesh-tone handling

*Column first appeared:* January 2005, *Electronic Publishing* magazine.

*Source of this file:* The author's draft as submitted to the magazine.

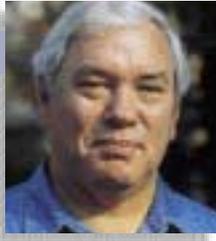
*Author's comment:* The column is notable for two reasons: first, as a demonstration of how eager (and gullible) readers can be; second, as the introduction of a major new technique. The gullibility refers to my opening demonstration that one can "sharpen" using only the Gaussian blur filter and layer blending modes. It was an excellent sharpen, but it took 23 steps, so the recipe had to be consigned to small type. I never imagined that people would be interested in something so convoluted. But if I had a dime for every person who's asked a question about that 23-step monstrosity I could have long since retired. The second reason is of more lasting value: I had introduced hiraloom sharpening in 2002 at the Seybold Conference and in my classes, but this was the very first mention of it in print anywhere as far as I know. Today, just about everybody uses some variant of it.

This archive, to be released over several years, collects the columns that Dan Margulis wrote under the *Makeready* title between 1993 and 2006. In some cases the columns appear as written; in others the archive contains revised versions that appeared in later books.

*Makeready* in principle could cover anything related to graphic arts production, but it is best known for its contributions to Photoshop technique, particularly in the field of color correction. In its final years, the column was appearing in six different magazines worldwide (two in the United States).

Dan Margulis teaches small-group master classes in color correction. Information is available at <http://www.ledet.com/margulis>, which also has a selection of other articles and chapters from Dan's books, and more than a hundred edited threads from Dan's Applied Color Theory list.

Copyright© 2005, 2010 Dan Margulis. All rights reserved.



## Life on the edge

Most images need sharpening, but one size doesn't fit all. Getting the best result depends on understanding that sharpening is really a form of blurring, that several strategies are possible, and that the key is the Radius

**Y**ou can't get far in the graphic arts without getting good at finding workarounds. The following is an unusual example. Figure B appears to the casual observer to be Figure A after Photoshop's Unsharp Mask filter has been applied. In fact, it's a workaround, a ludicrous but instructive workaround.

The challenge is to do this sharpening without the use of any filter other than the last one you might expect, Gaussian Blur, plus various layer blending modes.

Inasmuch as the full solution involves 23 steps, it's been consigned to small type in a small box. An abridged version, however, may help to explain why sharpening is usually necessary, and how it operates.

Unsharp masking is itself a workaround, a compensation, a fakeout method to create the illusion of focus. Almost every picture needs it, and not because photographers don't know how to focus their cameras.

Palm leaves have sharp, crisp edges—when we view them in person. On the printed page, it's a different

story. Each edge is far narrower than the halftone dots we print with, so regardless of how well focused the image might have been originally, there always seems to be a lack of focus.

Fortunately, painters, prepress people, and Photo-shoppers have known since at least the fifteenth century of a good workaround: darken one side of the transition, lighten the other, and let the viewer's imagination provide the knifelike edge.

To that end, we need something like Figure B, which outlines the leaves with a dark halo, and also darkens the vein lines of certain leaves. The whole thing gets multiplied into Figure A at a low opacity. Then, something that looks like an inverted version of Figure B, but in fact defines the parts of the edge that will be lightened, gets screened into the image to complete the "sharpening."

To produce Figure B, steer clear of Photoshop's clunky and inflexible Find Edges filter. Instead, make a

**Left to right, an original image; an intermediate file that used the Gaussian Blur filter to create the illusion of edges, which were then multiplied into the original as part of a preposterously complicated "sharpen" of Figure A.**





Greatly magnified sections of Figures A and C illustrate how sharpening algorithms strengthen edges.

duplicate layer, apply Gaussian Blur to the top, and change the layer mode to Lighten. That mode leaves the bottom layer intact, except where the blurred version on top is lighter.

Where light leaf butts dark background, the blur will fog the distinction, leaving the edge lighter on the background side and darker on the leaf side. But, since the blur layer is in Lighten mode, the darkening of the leaf is irrelevant. We now flatten this lighten-blurred composite, put it on top of another copy of the original image, this time in Difference mode, and flatten again. The new image is almost entirely black, which is what happens when the two Difference layers are almost identical. It isn't black in the areas where they aren't, to wit, the aforementioned background side of the edge. So, we use Auto Levels or some other method to make those difference areas stand out against the black, invert the whole file, and presto, Figure B.

If you'd like to save around 18 of the 23 steps enumerated below, you can use the filter whose name confounds many people who wonder why, if the purpose is to sharpen, it's called *Unsharp* mask.

By now you should be able to explain. Sharpening inevitably involves a defocused, or unsharp version of the file, which is then compared to the original with the idea of exaggerating the differences. In the days of drum scanners and, before them, process cameras, the unsharp versions were actually separate, deliberately out-of-focus exposures. Photoshop blurs the original to get nearly the same result.

Image technicians have known since before the birth of Thomas Knoll that this unsharp masking stuff normally should emphasize the darkening halos more than the ones that lighten the image. By the early 1980s, most drum scanners separated the lightening from the darkening aspects of USM, and by default emphasized darkening more. Amazingly, 20 years later, Photoshop still hasn't added this essential feature to what is arguably its most important filter.

Therefore, Figure C is in the spirit of drum-scanner sharpening. If you like it better than a straight application of the Photoshop filter, but don't own a drum scanner, the workaround is to make a three-layered file, sharpen the top two, set one to Darken and the other—at a lower opacity—to Lighten. The result is almost always technically better, but often the gain isn't worth the extra effort.

For images that *are* worth the extra effort (which are normally those being output at a large size—note how I've had to do some weird cropping of example images

so that their detail will be large enough in this magazine without having the article run 15 pages) there are at least seven other viable sharpening strategies, involving sharpening with different channels, through inverted luminosity masks, and the like. But the success of each one depends on our understanding and getting the correct answer to the following question:

How blurry should that unsharp version be?

### Shapen, not blurred

The pre-Photoshop answer to the foregoing question is, you get what you get. Drum scanners of the 1980s typically had 64K of RAM, not exactly conducive to flexibility in computing. Today's added firepower permits us to sharpen—if that's the word—images in ways that our ancestors never imagined. Of course, that won't help if we can't imagine them either. The point of this column to ensure that we can.

The last two columns dealt with a topic—lightening a grossly dark image—that seems entirely different from sharpening yet turns out to have a great deal in common with it. The key to Photoshop CS's powerful Shadow/Highlight command is control of the Blur field in its dialog box. If you grasp the significance of that setting, you'll understand why S/H sometimes seems to be sharpening the image as it lightens it, and you'll be able to extend the principle to USM.

Two sea pictures illustrate the possibilities. In each, the original is at left, a conventional sharpen at center, and a rad-

## How to Sharpen Using Only Gaussian Blur, in 23 Easy Steps

1. Make an additional copy of the image, naming it "For Darkening".
2. Add two duplicate layers to said copy.
3. Make a new copy of the layered image, calling it "For Lightening."
4. On the top layer of this new image, Filter>Gaussian Blur, 1.5 pixels.
5. Change layer mode on the second layer to "Darken".
6. Change layer mode on the third layer to "Difference".
7. Flatten image.
8. Make two duplicate layers. Set each to "Screen" Mode.
9. Flatten image and save.
10. Return to the image named "For Darkening".
11. On the second layer, Filter>Gaussian Blur, 2.0 pixels.
12. Change layer mode on the second layer to "Lighten".
13. Change layer mode on the top layer to "Difference."
14. Flatten image.
15. Invert the image.
16. Make two duplicate layers. Set each to "Multiply" mode.
17. Flatten image, save, and copy it to the clipboard.
18. Paste the "For Darkening" image on top of the original file.
19. Set layer mode to "Multiply."
20. Paste the "For Lightening" image on top of these two layers.
21. Set layer mode to "Screen."
22. Adjust the opacities of each layer if it seems advisable.
23. Flatten and save.

ically different approach at right. Both approaches work on the principle of putting an edge (or to use the preferred term, a *halo*) around transition areas, but the right-hand versions' halos are far wider, far softer.

Before beginning, two disclaimers. First, all images in this column are sharpened in RGB. Second, sharpening is notoriously subjective. If you think that some of them shouldn't be sharpened in RGB, you're right. If you think that I over- or under-sharpened any of them, you're wrong. But what I'm trying to do is show what different methods are likely to accomplish. Personal taste, and the colorspace being used, have no bearing on the message.

Conventional sharpening typically employs values of .8 to 2.0 in the Radius field of the Unsharp Mask filter. There is some dispute over how much to put into the Amount field, but values as high as the maximum, 500%, are by no means unheard of.

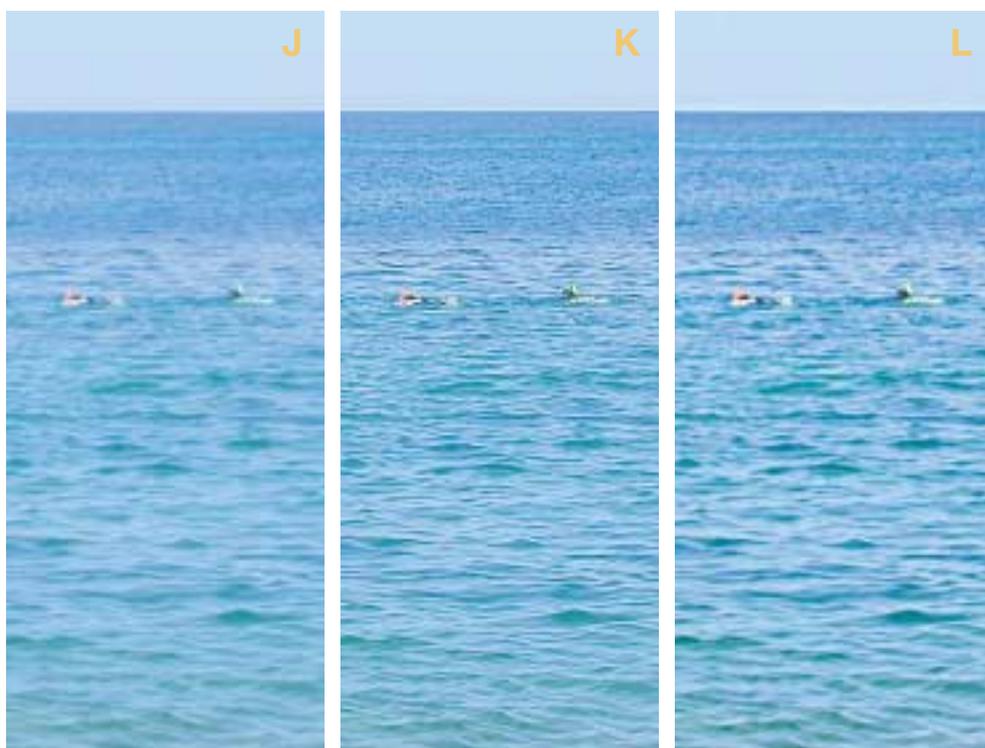
The alternative is a much higher Radius—say 10 to 25. Such an enormous blur will devastate the image unless accompanied by a drastic reduction in Amount—say, to 50% or, rarely, 100%. The result is very different from conventional sharpening, and should probably be called *shaping*, or perhaps *hiraloam* (high Radius, low Amount) instead. The result can be pleasing—sometimes. It isn't in Figure H. The big blur has sunk all detail in the light parts of the boats, with an enormous whitening halo. The masts have been improved somewhat, but still, one could easily argue that Figure F, the original, is better.

The conventional sharpening, Figure G, has no such problem. It has merrily found, and emphasized, lots and lots of individual edges, causing an appearance of better focus.

Figure J is a different story. There's no highlight to blow out, no shadow to plug, and nothing that looks like it needs to have a crisp edge. So, none of the reasons that *hiraloam* made a hash of Figure F are present. Figure K, the conventional sharpen, makes the



**Two sets of images, two styles of sharpening. Above, center, conventional sharpening works well in emphasizing well-defined edges. But below center, with no real edges apparent, it simply makes the water look grainy. The right-hand versions use a different strategy: a very high Radius, and very low Amount.**





**Fleshtone images are hard to sharpen without bringing out unwanted detail. Many people sharpen only the black channel, which contains little such detail. That was done in Figure X, but there was also a dose of RGB sharpening first.**

darkening of the background surrounding them, but there wouldn't be any of the gain in snap in the leaves themselves.

That's usually how it is. Images where shape-sharpening is unequivocally bet-

water look grainier, but I like Figure L better, because it's added shape: the rolling shadows in the waves are darker, the reflections in the water lighter, all without loss of detail.

If we tried this shaping method on the palm leaf picture of Figure A, it would lose again, but not as badly as in Figure H: the leaves would get lighter and there would be a pleasing

ter are rare, but they happen. If you have clients who furnish you that they been helpful enough to try to sharpen first, it would be well to know about how to do it, as conventional sharpening has big difficulties with files that have been damaged by hamhanded first efforts. Very grainy images, or older, damaged originals sometimes can't be sharpened conven-

# AD

tionally, but can accept a bit of hiraloom.

Sharpening fanatics look for the best of both worlds. Facial images, particularly those of women, are a real minefield, because any detail we bring out in the skin is likely to be something that the model wants us to see about as much as love handles or varicose veins. A conventional sharpen with a high Threshold will avoid this, at the unacceptable price of preventing sharpening of most hair.

A common solution is to sharpen only the black of CMYK, which usually has little or no facial detail but does define hair and eyes. But a hiraloom RGB (or CMY) method shapes faces effectively—if you can find the proper Radius.

To get this shape-sharpening effect, some people create a new layer, apply the High Pass filter, and set the layer to Overlay mode. It works, but it denies us the Threshold setting, and makes it much harder to find the proper Radius. If you want to sharpen in this style, use the USM filter, previewing at the absurdly high Amount of 500% to see exactly what's up.

Figure O was so sharpened, with a Radius of 5.0 and a Threshold of 0, creating something exceedingly unlikely to be favored by the model. Granted that a low Amount setting would soften this disaster considerably, there is no point in bringing out all these facial imperfections, when they can easily be eliminated by using a higher Threshold.

Figures P and Q, therefore, use a Threshold of 10 and Radii of 10.0 and 20.0, respectively. The blurring—er, Radius—value of Figure P doesn't quite wipe out the eyes. So, an edge develops, lightening the eye sockets and darkening the sides of the top of the nose.

In Figure Q, the blurring doesn't emphasize these features as much as it does the undesirable reflections off the nose and chin. The face lightens as a unit. So, 20 is too high; we pick something in between, and of course a much lower Amount. Figure N is the result, a sharpen at 70, 12, 10, followed by a light conventional sharpen of the black channel after conversion to CMYK.

**A warning to the wise**

This column barely hints at some of the sophisticated sharpening strategies for which a suitably deranged practitioner can conjure up workarounds blending and with combinations of different methods on different channels.

**When using high-radius sharpening, setting the Amount to 500% helps find the correct Threshold and Radius settings. Top to bottom: Radius 5, Threshold 0; Radius 10, Threshold 10; Radius 20, Threshold 10.**

If you'd rather stick to a single pass, remember that conventional sharpening became conventional for a reason. It would have been just as easy to design drum scanners to do hiraloom, but narrow-radius sharpening is best for most images. The trick, of course, is to recognize the ones where it isn't, and then be thankful that Photoshop lets us adjust both Radius and Amount.

Hiraloom can be tempting to those who are afraid to oversharpen images—it's harder to do something really objectionable to an image if you're using a low Amount, particularly if you do most sharpening in RGB.

On the other hand, there's a truism that people who sharpen in RGB (read: professional photographers) tend not to be disposed to (read: are so lily-livered that they won't) sharpen images enough. And, admittedly, it's very difficult to evaluate sharpening on a monitor, because the output usually appears softer than it does on screen.

The world will not come to an end if you oversharpen an image or two. You should ask yourself when the last time one of your clients complained to you that your work looked bizarrely edge-driven, or words to that effect. If it hasn't happened, you probably aren't sharpening enough. Clients pay good money for the privilege of griping about our work. If we're doing our job right, they should be equally likely to complain about oversharpening as that the image lacks snap; equally likely too colorful and not colorful enough.

For most of us, oversharpening isn't the typical complaint, which is a very good argument for learning how the Radius does its thing. An overly soft image once in a while is no big deal, but once you get a reputation for producing dull, soft-looking output, it may be a sharpening problem for which there isn't any workaround.

.....  
 Contributing editor **Dan Margulis** (DMargulis@aol.com) is author of *Professional Photoshop, Fourth Edition*. For information on his color-correction tutorials in Atlanta, Chicago, and San Diego, call Sterling Ledet & Associates at 877-819-2665. To join Dan's on-line color discussion group, [www.ledet.com/margulis](http://www.ledet.com/margulis).

