

Image Sharpening-Make it Really Sharp

Jeff Schewe

A variety of tools can sharpen images. You can sharpen in your scanner or digital camera, which is rarely ideal and often destructive. In Photoshop you can simply use the USM filter. While it's a critical core tool, it's widely misused. There are other recipes-most of them obscure and lengthy-for doing different kinds of sharpening within Photoshop. There are even 3rd party plugins that strive to offer a wide array—some might say a confusing array—of sharpening methods. Some users dodge the question and try to leave sharpening to others—preferring to hand off conservatively under-sharpened images that really are too soft-but not badly stepped on to others in the hopes that it'll no longer be THEIR problem and will magically turn into somebody else's. The problem with this strategy is that it's just as likely to back fire as it is to produce optimum sharpening in prints or reproduction. The problem is that nobody really has developed a consistent and optimum “Sharpening Workflow”. Well, till now.

Some Sharp Fundamentals

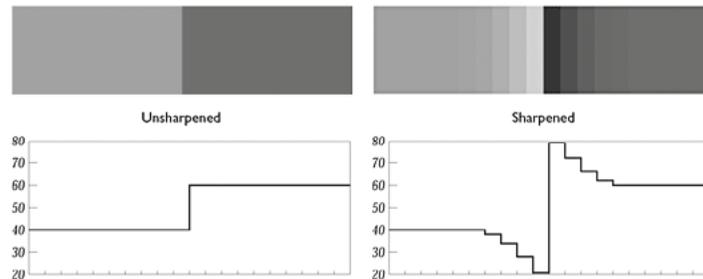
The first aspect you must understand is why we need to sharpen digital images at all. Shouldn't scanners or digital cameras be designed to produce optimally sharpened digital images? In a future universe with artificial intelligence, maybe that could happen. But until we transport through time into the future, we're stuck with the basic fact that all images are not created equal and the technology to digitize images is basically “soft”.

When we capture images, we use imperfect lenses (though some are more imperfect than others) to focus the photons onto a noisy capture medium. Film has grain, digital sensors have a fixed grid of pixels, and each softens small details. Most one-shot digital captures are further complicated by the fact that each pixel in the sensor records *only* red, green, or blue information—the full color information is then created by interpolating the missing color channels, and this process is prone to producing strange artifacts—stair-stepping or stray oddly-colored pixels—that most camera vendors avoid by introducing some blurring, either optically or during the software processing, losing sharpness. When we scan film, the grain in the film interacts with the scanners' fixed pixel grid, introducing noise artifacts that mask detail AND reduce our ability to apply global sharpening.

Then, when we reproduce images for output, the various processes we use to do so, whether film recorder, inkjet printer, or offset press, all introduce some additional softness in the process of turning pixels into dots of ink or dye. Each output process has a limit to the finest detail it can reproduce, so we have to optimize the image sharpness differently for each output process

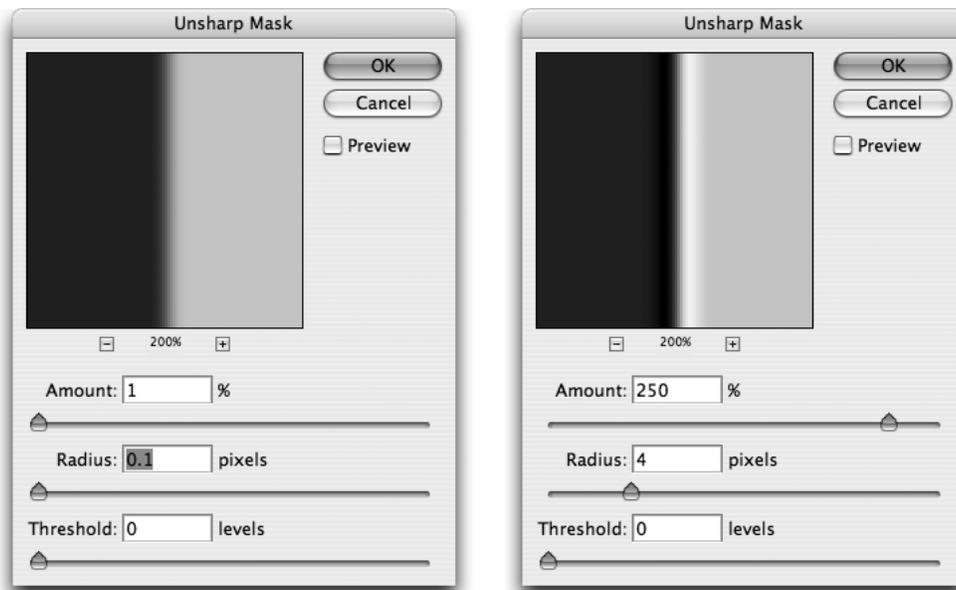
To increase the apparent sharpness or acutance, you can resort to algorithmic processing to put a sharper “edge on”. Unsharp Mask in Photoshop (USM) is one such process. I could give you the math (although I personally don't really understand it) that describes the process-but I won't. What USM basically does is increase the image contrast between light and dark pixels. The idea is to process the light side of edges to be a bit lighter and the dark side to be a bit darker. This will indeed increase the apparent sharpness, but it introduces additional problems such as

increased film grain or camera noise, or in extreme cases blown out detail and a real scrunchy look. A big part of the problem is that the USM filter doesn't know whether the contrast between adjacent pixels represents an actual edge or a noise artifact such as film grain. The standard USM approach of selecting an Amount (the strength of the process), the Radius (the edge radius where the process produces the contrast increase) and the Threshold or the number of individual tone levels where the process will be applied. These three controls offer very little in the way of fine-tuning the USM process-and fine-tuning, on an image by image basis is where you can really separate the "Men" from the "Little Boys".



In this figure, the effects of a sharpening algorithm are shown. The image and graph depict an edge transition from 40 to 60 percent. Each tick mark across the bottom of the graph represents a column of pixels. After sharpening, the transition is accentuated-darker on the dark side and lighter on the light side. This creates a halo that increases the apparent sharpness. Image courtesy of Bruce Fraser from *Real World Photoshop*.

The single biggest limitation to the USM filter is that it's a rather blunt edged tool for such subtle image surgery. It's all too easy to over-sharpen. The other limitation is that when used "globally" over the entire image, it will almost certainly hurt some areas of the image while helping other areas. . .this limitation requires some sort of local discrimination and intelligence. You should only apply sharpening to those areas that need it-and this is the primary reason that global applications of sharpening applied by a scanner or digital camera pretty much suck. There's no finesse when you apply sharpening equally to all areas of an image.



A soft edge on the left can increase in apparent sharpness after applying USM. Note that the lighter edge is lightened while the darker edged is darkened thus increasing the apparent sharpness

But the finesse needs to go further. The ideal sharpening will be directed towards only those areas that benefit from sharpening such as edges, and masked off those areas such as film grain or camera noise. You also need to make sure you control the delicate highlight texture and cut down on the amount of over-edging or haloing that sharpening can produce.

Another key factor is the nature and the content of the image itself. The image can tell you a lot about the sharpening it needs. Images with a lot of high frequency textural detail such as a landscape with sharp rocks and trees demand an entirely different approach to those with low frequency detail such as a person's face. You want to sharpen the eyes and the lips, but you need to either eliminate any skin sharpening or actually introduce a level of smoothing. Of course, most images are a combination of types of texture and detail, which further proves that no single "one size fits all" global sharpening will ever be ideal. This stuff isn't rocket science. . .rocket science is a lot easier. . .either you land on the planet or you don't. Having the skills and finesse to produce a perfectly sharpened image takes training, practice and experience-and darn good eyes.



The face on the left is an example of a low frequency image while the fern on the right is a high frequency image.

And speaking of eyes, just looking at an image really isn't good enough. Since CRT or LCD monitors can never match the resolution of either a fine ink jet or halftone print, you are always placed in the precarious position of predictive sharpening. Predictive evaluation takes good monitors that are well calibrated and profiled (since color and gamma can indeed effect the appearance of sharpness) as well as an ideal environment to be viewing the monitor's image on. Many people will tell you that in order to judge how to sharpen an image, you need to evaluate the sharpening effect at a 100% screen zoom ration. Well, yes, in order to see the actual effect any sharpening will have on your PIXELS you need to view at 100%, but in the case of the resolution of monitors, this will over predict the amount of sharpening. To more accurately judge the amount of sharpening to apply, you really need to "soft proof" the final output. So, if you are providing a 300PPI image intended for halftone reproduction at a 150 line screen, you are providing 4 pixels for every halftone dot. To soft proof the amount of sharpening you need to

apply, you really need to judge the amount at a screen zoom ratio of 50%. Thus, four image pixels will produce a single monitor pixel—a similar dithering to the halftone screening.

Since the zoom ratios within Photoshop that produce nice pixel dithering are limited, I suggest also using a 50% ratio for ink jet and contone printing as well. It's very important that you avoid trying to judge sharpness at the "odd" pixel ratios—33.3%, 66.7%, and so on, because Photoshop applies a hefty amount of anti-aliasing to those views, making it impossible to predict the actual sharpening with any degree of certainty. Until such time as monitors increase their pixel densities, we're kind of stuck with these limitations.

Another issue affecting the ability to judge the amounts of sharpening needed is your monitor's actual resolution. While your high end 21" monitor may be capable of viewing 1600x1200 image pixels on the screen, those are NOT real pixels. They are produced by aperture grills and masks that dither the image to produce the 1600x1200 resolution.

One of the real benefits of using digital LCD monitors such as the Apple Cinema Display is the fact that you are actually seeing REAL pixels. . .not the result of a dither. So the screens appear far sharper and the ability to soft proof sharpening IS SIGNIFIGANTLY IMPROVED.

Ok, so much for the primer on all the reasons why, and some info on how and some tips on viewing. . .the \$64,000 question still is, how to implement an entire sharpening workflow into your processing?

Well, you COULD just buy PhotoKit Sharpener, a product from some of the usual suspects at Pixel Genius. Bruce Fraser, Martin Evening, Seth Resnick, Andrew Rodeny, Jeff Schewe (that's me!) and our engineer Mike Skurski have designed a product around the philosophy of a managed sharpening workflow. But that's like giving a fish to a hungry man, it may give him a meal of sushi but it won't teach him how to catch fish to feed himself (let alone how to make sushi). So instead, we'll tell you how our sharpening workflow actually works.

First off, you must develop a series of recipes for each image source you may use. If you generally shoot the same type of film and always use the same scanner, this will be pretty easy. Same deal if you pretty much always shoot with the same digital camera—just understand that within these broad categories, there's still a need for different flavors. Chromes need one kind of sharpening and negs require another. Large format needs another while 35mm yet a different flavor. Even with digital capture, you may need a different routine for raw and jpg and different routines for each ISO you shoot at. The input or "Capture Sharpening" is the first leg of the sharpening workflow. The aim is to introduce the correct amount of pre-sharpening that regains the sharpness lost during digitization while being very careful to not do any harm to the image. You need to work the edges and consider the nature of the source to produce the correct sharpening process. Don't worry, I'll offer some brilliant techniques, but I wish to evangelize the workflow so you understand the differences each stage needs.

Even if you nail the exact correct amount of sharpening suitable for the pre-sharpening stage, you will still need to do localized sharpening based upon the needs of the actual image. This is where sharpening becomes a creative tool. The ideal method is either to use selections to mask

the sharpening, or use what I prefer-sharpening brushes. Didn't know Photoshop had those? The point of the "Creative Sharpening" is to add and enhance the image. The capture stage regains the sharpness lost but the creative stage is where you can really turn a sow's ear into a better sow's ear. Seriously folks, if you miss focus, have camera shake or subject movement in the image, there's only so much you can expect to do about it. I mean, the special effects on TV shows like CSI are neat, but if an image is soft, blurry or has movement, you can NOT magically turn the image sharp when your boss says "can you enhance that?". That said, it's remarkable just how the artistic use of localized sharpening can substantially improve an image. Also, don't forget the old trick, if you want something to look sharper, you make it appear sharper by making everything around it softer.

So, even though you've jumped through the hoops of capture and creative sharpening. . .you still aren't there yet-sorry.

The last stage of the sharpening workflow is "Output Sharpening". This stage is where you sharpen the image-*after final sizing*-based entirely upon the type of output and media you may be using. If you're printing to an Epson ink jet, you'll need more sharpening for matte media than for glossy media. If your image resolution is 240PPI, you'll need a different sharpening than if you are printing an image at 360PPI. Same deal for halftone output-a 175 line screen at an image quality of 2X (twice the PPI rez as line screen) on coated stock will need to be sharpened differently than an image printed on 133 line screen uncoated stock. Also, don't forget the web or multimedia. If you want to sharpen a 400 pixel image, it'll be different than a 100 pixel thumbnail.

I'll admit, this may all sound overwhelming. It is. This is where the real sharpening pro can be distinguished from the rookie. But, if you understand the principles, it will go a LONG way to making the process understandable. One method of thinking about this is to compare sharpening to a color management workflow (oh goodie, comparing one unfathomable workflow to another-yeah, that helps).

Actually, it's not that bad. Think of your capture sharpening as the conversion from the scanner or camera into a color working space-except in this case it's sharpening. Then once you get an image into Photoshop, you would logically do those local corrections needed by the image-but instead of doing color correction you're doing sharpening enhancements. Then when you have the ideal optimized image, the final stage is to sharpen for the output.

When you break the segments down to their logical steps and you employ good soft proofing for predictive sharpening, the whole messy process begins to take a manageable shape. You just need to do the proper incremental sharpening as you proceed from stage to stage.

So, assuming you understand the concept, here are some actual techniques for various sharpening routines. All of the following recipes are totally friendly for both Photoshop 6 and 7 and will work in Photoshop CS even in 16 bit and require no additional tools.

The first step is to learn to do “Edge Sharpening”. This is the process of producing an edge mask for applying a certain level of sharpening through. There are a variety of different formulas but this is a pretty simple process.



The image on the left is the full frame while the one on the right is the detail I'll be using so show the steps.

The following steps are for creating an edge mask to use for sharpening edges. The creation of the mask can have many variables and various formulas, but this is a simple recipe.

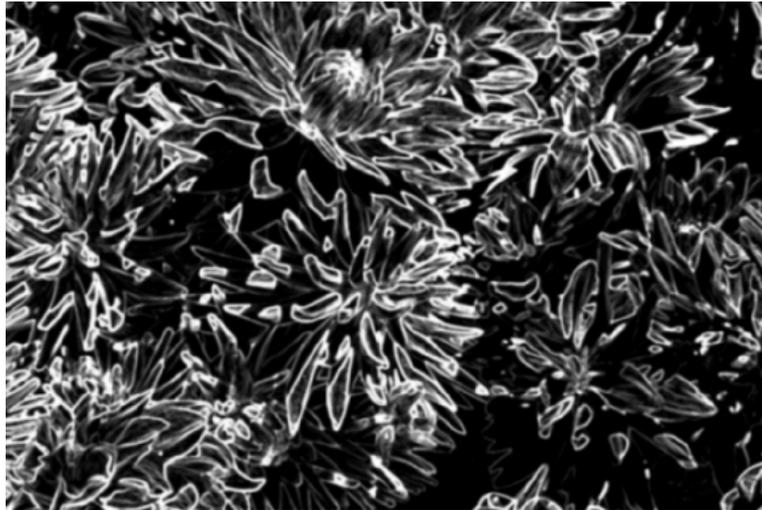
1. In your image, duplicate the Green channel-this channel has the largest component of luminosity information.
2. Name this channel “Edge Mask”
3. Run the Photoshop Filter named Find Edges (under the Stylize menu). Run the Median Filter to expand the edges-use 1-3 depending on file resolution and edge width (under the Noise menu).



A copy of the green channel after running the “Find Edges” filter and the levels command to adjust the thickness of the edges.

4. Use the Levels command to darken the edges. Run Gaussian Blur to soften the edges- about 1.5-3 pixels depending on file resolution.

5. Invert the channel colors (under Image>Adjustments>Invert).



This is the resulting “Edge Mask” to load as a selection and apply edge sharpening only to the lighter areas of the mask.

To use the mask, the ideal way is to duplicate your image that you wish to sharpen to a separate layer (command J) and set the blending mode to Luminosity. Using luminosity is similar to converting your image to Lab and sharpening the lightness channel but without introducing the color conversion data loss. Set the layer’s opacity to in intermediate opacity-I like to use 66% to allow upwards and downwards opacity adjustments. Then load your edge mask as a selection and run the Unsharp Mask filter. Because you are already constraining the width that the USM filter will work on, you can be more liberal with your settings. The threshold is not important so leave it at zero.



Applying the USM filter through the edge mask on a duplicated layer set to Luminosity.

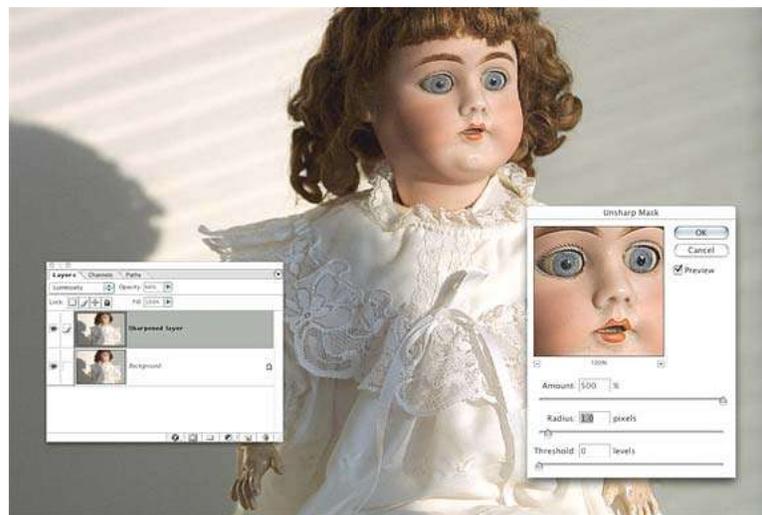
The edge sharpening routine can be recorded as an action for batching to multiple images. This edge sharpening routine is perfect for capture sharpening.

When it comes to localized creative sharpening, the basic concept is to apply a certain level of sharpening and then localize the effects by selective blending opacities. This process is real easy if you use layer masks.

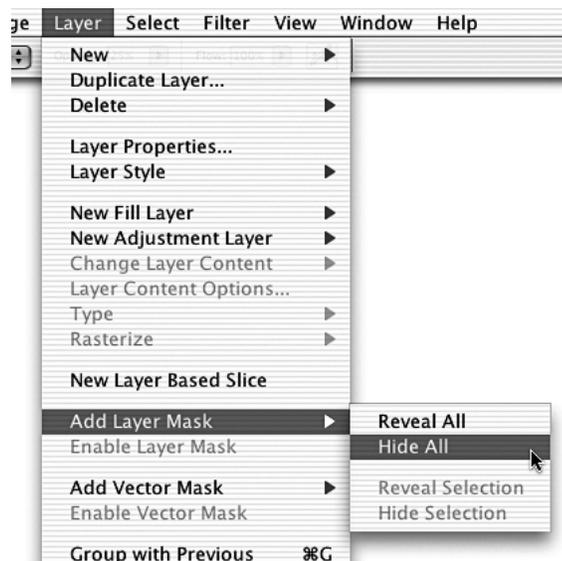
Step-By-Step: Sharpening Brush

The purpose of the sharpening brush is to apply sharpening localized sharpening only to those areas that you wish.

1. As in the other sharpening methods, duplicate the image to a new layer.
2. Set to Luminosity blend and an opacity of about 66%.



3. Run the Unsharp Mask filter-you can run it pretty strongly since you'll be only using a low amount of the sharpening.
4. Add a Layer Mask using the Hide All Option.



5. Set your foreground color to white, select a large soft brush with a lowered opacity (I suggest about 20%) and simply paint into the areas where you wish enhanced sharpening. If you goof up, simply switch your paint color to black and you can paint away the sharpening areas.



This localized sharpening is far superior to merely slapping on sharpening over the entire image. For creative sharpening, this is the ideal method of fine tuning the overall sharpness of an image.

The final stage of the sharpening workflow is output sharpening. Now, to be honest, you really have to know some stuff before you can do the final stage. You **NEED** to know the final size and more importantly the final resolution of the image. You also need to adjust for the type of output you are doing. Generally, halftone sharpening can be done using a certain level of classic USM in Photoshop. But the trick here is to only go so far, and no farther.

Output Sharpening

Output sharpening for optimum ink jet printing can be a bit tricky in that a single sharpening method really isn't ideal. I like to combine a small amount of USM with another sharpening process called "High Pass Sharpening". In this process, you actually use the edges of the image layered and set to Overlay blend mode and then run a filter called High-Pass. In overlay mode, you are actually using a procedural blend where the lights get lighter (Screen) and the darks get darker (Multiply). The High Pass filter is a variant of an edge finding algorithm that preserves the edges and turns everything else into a middle gray-approx 128 RGB. In overlay blend, middle gray does nothing. So the net effect is to lighten the lighter boundaries of the edges while darkening the edges. This process produces an optimum edge halving while allowing for a more precise control over the sharpening by controlling the opacity and blending options.

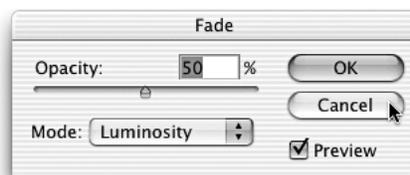
Step-By-Step: High Pass/USM Combo Sharpening

One of the ideal ways to sharpen for final output is by using a combination of what's called "High Pass Sharpening" as well as a small dose of Unsharp Mask sharpening. The critical factor in this process is that the image must be sized to the final output dimensions and resolutions before the application.

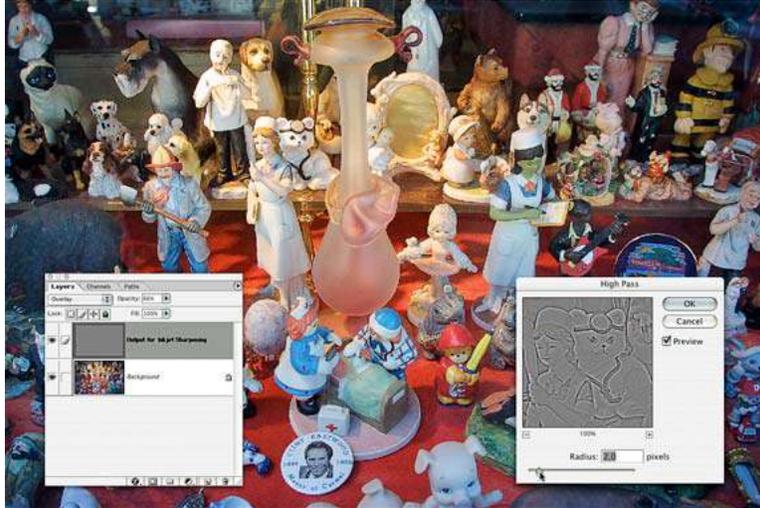
1. Duplicate the image layer for sharpening and leave the blend mode at Normal and set the Opacity to 66%.
2. Run the Unsharp Mask filter at a moderate amount making sure to put in a threshold (3-6 depending on the grain/noise of your image) to avoid over sharpening the grain or noise)



3. Immediately go to Edit>Fade and set to 50% (or your own flavor) and set to luminosity blend-we're doing the luminosity blend here instead of the layer blend because we'll be changing the layer's blend mode shortly.



4. Now change the blend mode to Overlay-don't worry if the image takes on a harsh contrast, the next step fixes this contrast effect.



5. Go to Filter>Other and select High Pass. The exact amount needed will be dependant upon the exact resolution of your file and it's intended media-glossy needs less and matt paper needs a bit more. A good starting point is 2 pixels. This will localize the overlay contrast increase to just the edges in a manner very similar to making an edge mask.

By combining a small amount of USM with the edge sharpening offered by using High Pass on an Overlay mode, you can further refine the sharpening needed by your output device. After running the sharpening, you can adjust the overall sharpening by setting your screen zoom to 50% and adjusting the opacity slider. Also, don't forget the advantages of recording commonly used steps as actions so it is easy to apply to your images.



So, what do you think? Does the concept of an image “Sharpening Workflow” make sense to you? It does for me. Taking a workflow approach allows you to concentrate on a stage by stage approach to sharpening. It allows you to limit each stage to only those goals needed while allowing for a build up of the total sharpening needed for your final output. The addition of producing the sharpening effects on layers also allows you to modify them to taste on an image by image basis. While breaking the workflow down to three distinct stages may seem like an additional work burden, aside for the creative sharpening which requires human intervention, all of the capture and output sharpening can be automated by creating actions and batching them on a large number of images at a time. Is it easy? Nope, I never promised you EASY, but at least the workflow concept makes the process fathomable-and that is a big step forward.

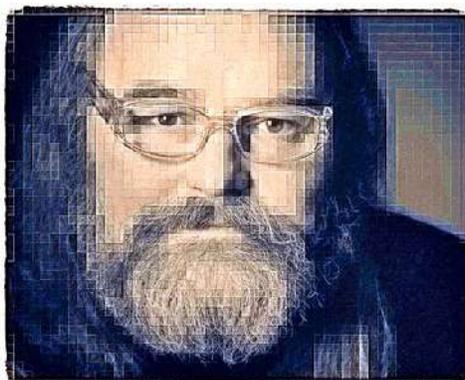
Additional Sharpening Information:

Creative Pro article by Bruce Fraser about a “Sharpening Workflow” can be found at: <http://www.creativepro.com/story/feature/20357.html>

An additional article about sharpening by Bruce Fraser can also be found at: <http://www.creativepro.com/story/feature/11242.html>

Information about PhotoKit Sharpener can be found at: <http://www.pixelgenius.com/sharpener/index.html>

Sample Sharpening workflows by Seth Resnick and Bruce Fraser can be found at: <http://www.pixelgenius.com/sharpener/workflows.html>



Jeff Schewe is a Photoshop Guru’s Guru. He’s on the inside of the development and testing of Photoshop and has helped guide and direct many features since Photoshop 4.0. Short of Thomas Knoll and Mark Hamburg, there’s probably not many people who knows Photoshop like Jeff. An award winning advertising photographer of over twenty years, he’s a grizzled veteran of creating the impossible image often under impossible conditions. He has developed tricks and techniques even the Photoshop engineers didn’t know could be done. As an indication of his skills and knowledge, he has been named a Canon “Explorer of Light “ as well as an “Epson Stylist Pro”. He is a past Apple “Master of the Medium” and a former member of the Apple Customer Advisory Board. Additional information is available at: <http://www.schewephoto.com>.