Digital Sharpening I

February 2004

We want our photos to be sharp. We spend our hard-earned money on the very best lenses, the latest cameras and the highest resolution printers. It is well known that we need to sharpen our digital images - every book on digital imaging devotes a couple of pages to sharpening. However, when it comes to the nitty gritty of sharpening a particular image, even experienced photographers are often lost in the maze of different methods and options. Without a good understanding of sharpening techniques the results will be average at best, and often poor.

The goal of this series is to give us a clearer picture on what we mean by sharpness and to shed some light on the mysteries of digital sharpening.

Sharp, sharper, sharpest

When looking at an image, our eyes are scanning for edges. Visual information is broken down into 'edge' and 'not edge'. Not every, but most good images have a strong edge definition.

Sharpness is made up of two factors: **resolution** and **contrast**, which are interrelated. Resolution is measured by the ability to resolve a certain number of lines per millimetre (l/mm) at a given contrast. Fuji's Velvia for example, has the resolving power of 80 l/mm at a chart contrast of 1.6:1 and 160 l/mm at a contrast (or luminance) ratio of 1000:1.

For the lens designer, contrast and resolution are in conflict: increase one and the other one goes down. This has historically been a battleground between Zeiss, who designed their lenses for maximum resolution and Leica, who tended to favour higher contrast. (For the mathematically inclined reader I recommend to have a look at the modulation transfer function (MTF), which allows evaluating the sharpness of a lens with a simple diagram. Norman Koren has a very good introduction to MTF on his website, www.normankoren.com).

There is a third factor contributing to our perception of sharpness: film grain or its digital counterpart, noise. Often an image with moderately coarse, but sharp grain can appear as sharp, or even sharper than a very fine-grained image. The same applies to digital cameras. I have seen images, especially images without any fine details, taken at 100 ISO and totally free of digital noise, but they appeared less sharp than a 400 ISO frame from the same camera.

Why do we need to sharpen digital images?

Softness is introduced every time we cross from the real, continuous, analogue world into the digital world. Scanners and digital cameras, all need to transform continuous gradations of tone and colours to points on a regular sampling grid. Details finer than the sampling frequency get 'averaged' over the pixel grid, softening the overall appearance. Most digital cameras use the Bayer Pattern CCD, where each CCD records only one of the three primary colours and the other two colours need to be interpolated, which again contributes to softness. Output through press halftones and inkjet dither patterns further introduces softness, when image pixels are converted to dots of ink or toner.

Whatever the source, a simple flatbed, the most sophisticated drum scanner or a digital camera, the image needs to be sharpened. Photoshop gives us a number of options. Forget about Sharpen, Sharpen More, Sharpen Edges and the Sharpening tool. They may be useful for creative effects, but will wreck your images very quickly if you try to use them to compensate for the softness introduced during either acquisition or output.

Unsharp Mask (USM) is the only readily available filter to do the trick. As we will see later, there are a couple of other ways to sharpen an image without Unsharp Mask, but it is and remains the most commonly used method. Let's have a closer look.

Unsharp Masking

This may sound like the last thing you would want to do to your images to make them look sharper. In my darkroom days (it feels like the Dark Ages now) there was only one way of getting the print to look sharper, namely to use a 'harder' paper. In the commercial printing world, the film was sandwiched in the enlarger with a slightly out-of-focus duplicate negative - the unsharp mask - and the exposure time for printing was approximately doubled. This made the light side of an edge to print lighter and the dark side to print darker. The resulting halo along edges and around objects made the image to look sharper to the eye.

That is how the Unsharp Mask filter got its name, and it is supposed to work similarly to its darkroom predecessor.

Here my curiosity got the better of me and I had to simulate this procedure in Photoshop. I took the image of a galaxy, an image which doesn't lend itself very well to Photoshop's USM filter, because it hasn't any normal, strong edges. I duplicated the layer and applied a fairly generous dose of Gaussian Blur. Then I reduced lightness and contrast of the blurred layer, put a copy of the original image on top of my 'unsharp mask' and set the blending mode to Difference. Finally I pushed the white point in Levels to the left, to simulate the increased exposure time. The result was impressive and very similar to, but not quite the same as using the USM filter with a heavy radius setting.

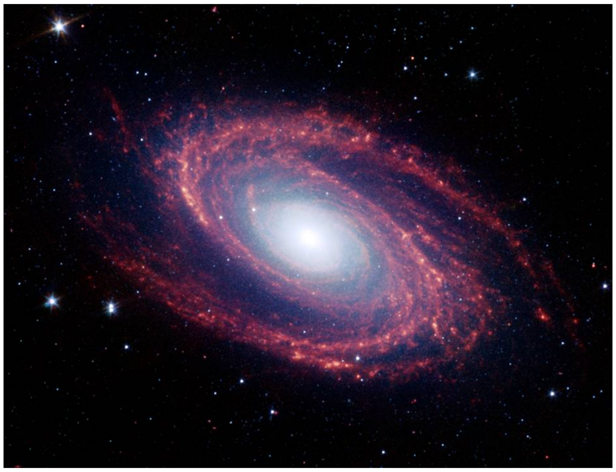

Original: Messier 81 Galaxy (Spitzer Space Telescope) - NASA

The 'Unsharp Mask'

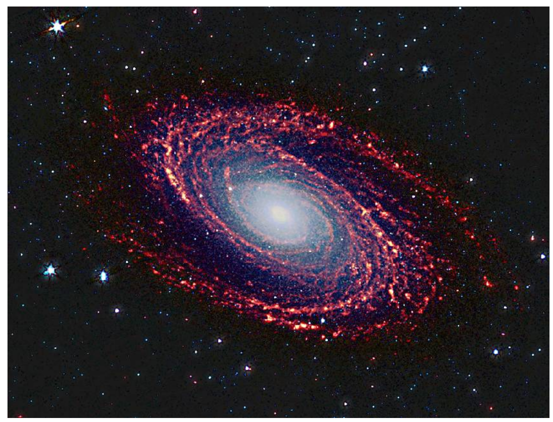

The sharpened galaxy

Let's now turn our attention to the USM filter, which you will find in every image-editing program.

How the Unsharp Mask Filter works

The USM filter works by going through the whole image, pixel by pixel and comparing them to their neighbours. If there is a certain difference in contrast between adjacent pixels, the filter will exaggerate the contrast according to the parameters you set. Here is where the fun starts, because the filter gives us three sliders - Amount, Radius, Threshold. They all have an effect, but finding the optimum setting seems a hit and miss process, if one doesn't know what the sliders are actually doing.

I illustrate the principle of the USM filter on the opposite page with the letter 'A' in dark-grey on a light-grey background. The close-up shows how the filter increases the contrast by lightening the light pixels and darkening the dark pixels along the edges.

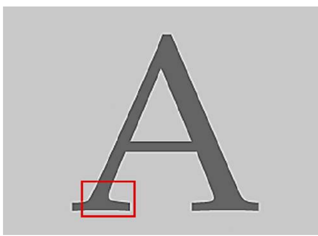

Not sharpened

With USM Filter

Here we can already see possible problems coming up: If the affected areas are too wide, then the sharpening effect will be visible from a normal viewing distance as a distracting halo. Secondly, details in the highlights could be blown out and shadow details can get lost in the process. Thirdly, the filter doesn't know if there is an edge that we want to accentuate, or if it is film grain, camera noise or a pimple on the face of our model, which we don't want to sharpen.

How to set the Controls

When applying sharpening make sure that you view the image at 100% magnification to see the effect on fine image details, but also keep an eye on your image at 50% - it will give you a better idea how the image will look like when printed.

The best starting point is the second of the three sliders - **Radius**. This determines the width of the halo created by the filter along edges. The wider the halo, the more pronounced the sharpening. But we need to avoid an unnaturally oversharpened look. The setting depends on the file size (the number of pixels in the image) and the image content. A good starting point for a 10-40 MB file with fine image details will be a setting of '1'. Use a high 'Amount' setting, say 300% and a 'Threshold' of 0. The Radius can be increased to 1.4-2.0 for larger files and images without fine details. Small files for monitor display will need a much smaller setting, around 0.4-0.7.

The next control is the **Amount**. This is the volume control of unsharp masking. It determines the intensity of the halo. At the high end, pixels will be forced into either pure white, solid black or to the maximum amount of a given colour. Start with a high setting, say 400%, and work your way down until the image starts to look good.

The **Threshold** setting tells the filter not to sharpen image tones which are similar. For example, a Threshold setting of 5 will ignore all tones that are within 5 level values of each other. A setting of 3 is a good starting point for the average image. A figure of 6-8 will protect noisy shadow areas, skin pores and other small blemishes from being accentuated.

When to apply Sharpening

Sharpening is a destructive process that permanently alters the pixel contrast values. Once an image has been oversharpened, the damage cannot be undone. It makes sense to sharpen on a separate layer, or to keep a copy of the unsharpened version.

Sharpening can be applied in stages and it is recommended to use a three-step approach: A first pass sharpening to correct for the softness introduced by the capture device, a second step during editing for creative sharpening applied only to important areas of the image, and as a last step, a final sharpening targeted for the output device.

Softness is introduced every time we resize the image. It is therefore important to do the final sharpening **after** the image has been sized correctly for print or monitor display. Sharpening should be the last step in our editing chain, just before CMYK conversion.

Advanced Sharpening Techniques

There are various tips and tricks to get the best out of the USM filter. We can sharpen individual colour channels, or switch to Lab mode. Sharpening can be targeted to edges only with a mask ('smart sharpening' as shown in the B&W image below).

There are also other ways of sharpening without the USM filter. I will discuss these refinements in next month's issue.

Demonstration of 'Smart Sharpening'

(To be explained in Digital Sharpening II)

Autumn in Vienna - FP4, ISO 400

No sharpening

'Smart Sharpening'

Normal USM Filter