INTRODUCTION TO GRAPHICS

LIGHTING GRADIENTS

Information Sheet No. **PS824**

Shading or brightness gradients in images may be visually ignored because the human eye doesn't particularly notice gradual changes in brightness. But we do tend to overlook the darker areas of photos, so one reason to correct for shading is to give all parts of the scene the same average brightness. When we look at the dark region of a scene we can adapt our vision to bring out the detail. But the dynamic range of a digital photograph is not great enough to retain the subtle details in extreme light and dark areas. Correcting for shading or brightness gradients can allow enough dynamic range for enough local contrast to make detail visible everywhere.

In a studio setting, the placement of lights controls the levels of illumination, and it is possible to arrange them so that all of the interesting parts of the scene have adequate light intensity. It may even be desirable to have the background areas darker, to focus attention onto the foreground. If it isn't practical to place all of the lights needed to illuminate the subject perfectly, acquiring multiple images with different lighting and combining them may provide an effective solution. Particularly in macro photography, it can be very difficult to arrange uniform lighting. Ring lights around the lens, or an array of floods on both sides, can solve the problem in many situations. But sometimes, even with the best practical effort, the recorded image will have a brightness gradient. Also, optics can contribute to the problem. In many situations a mismatch between lens, apertures and chip can produce vignetting, in which less light reaches the edges and corners of the chip and the image is darker away from the center. Many of the less expensive consumer cameras use extremely low cost lenses that often suffer from vignetting problems, especially if limited capability optics are asked to provide a zoom capability (these lenses also have out-of-focus and distortion problems away from the center of the image).

For whatever reason vignetting or side-to-side brightness gradients may arise, if they can't be eliminated before the image is taken it may be necessary to correct for them afterwards. The most straightforward way to accomplish this is to take a second image with the same lighting and camera settings, but without the subject being photographed. In place of the subject, a uniform grey card is used. The image of a standard "18% grey" test card will record the lighting variations on the subject and can be used to level the image brightness and remove vignetting or gradients.

The name "18% grey" may be misleading–because of the logarithmic response of the eye, film cameras, and many digital cameras, this card will actually be seen as a medium grey and should ideally correspond to stored values halfway between white and black (in other words, about 128 on the 0..255 brightness scale). Since it is a neutral grey, it should have equal intensities for red, green and blue. Some copy stands incorporate a grey card in the base, so simply removing the macro object being photographed permits recording the grey background image.

The first illustration shows an example of rather severe non-uniform lighting (one of the floodlights on the copy stand was turned off). Removing the necklace and capturing a second image with the same camera settings produced a background picture that recorded the texture in the cloth used as a background. To eliminate this from the background shading image, a large Gaussian blur was applied.

Background removal using a measured background image: a) original image with non-uniform illumination; b) background with necklace removed; c) image a divided by a blurred copy of image c and adjusted to full contrast range.

(It sometimes occurs to people who dimly recall some of their high school algebra that there might be an alternative way to accomplish the division by using the Photoshop Image->Adjust->Invert command and then the multiply option in either the Image->Calculations dialog or in the layer blending mode. They are thinking that dividing by X is the same thing as multiplying by the inverse of X. But, alas, that doesn't apply here. Inverting an image doesn't calculate the inverse in the sense of 1/X but rather inverts the image contrast by subtracting the pixel values from 255. Sorry about that...) Because the background image has a brightness range that covers about 25% of the full 0..255 dynamic range of the original image, the resulting image has lost some of its brightness resolution, but by expanding the remaining 75% to cover the full range (using the Adjustments -> Levels function) the image is still quite satisfactory. If the background variation becomes too large, for instance 50% of the dynamic range, it creates problems because the foreground objects have too little contrast remaining, and some regions may be clipped to full white or black.

Generating a Background

Obviously, the grey card method is a lot harder to implement in an outdoor scene. It may be possible to locate features in various parts of the image that should have the same brightness and color (for instance, the shirts on team members in a group picture taken with lighting from the sides that makes the center darker). If several such locations can be found, they can be used to construct a background image.p> In the example shown below, the brightness values on the shirts at the left and right sides of the image were measured and used with the manual gradient tool to construct a ramp. This was done in grey so that the colors in the original image would not be altered. Then the background ramp was divided into the original image, pixel by pixel, to remove the shading. This reduces the dynamic range of the image, but after increasing the contrast top cover the full display range using the Adjustments > Levels function, the result is an image that does not show the shading.

The built-in Photoshop calculations routine offers subtraction but not division. For a relatively small amount of shading, as in this picture, subtraction of the gradient produces a nearly identical result and can be used satisfactorily. When the extent of shading is greater, it becomes important to construct the new image correctly as the ratio of the original image to the shading. Plug-ins such as The Image Processing Tool Kit and Fovea Pro offer ratioing with optional automatic scaling of the result.

Removing shading from an image: a) original image; b) linear ramp constructed with the manual gradient tool, using brightness values measured on shirts at the left and right sides; c) image a divided by image b (pixel by pixel), with levels adjustment to increase contrast.

More complex backgrounds (using polynomials instead of linear ramps) can be constructed automatically by selecting multiple regions in the image and using the Background Fitting and Background Removal plug-ins in Fovea Pro. The next illustration shows using the Select > Color Range dialog to select all of the yellow areas on the boys shirts in the original image shown above. The automatic fitting of a background function to these pixel values and its removal produces the corrected image shown.

Automatic shading correction: a) selecting the yellow shirts with the Select > Color Range function; b) removal of the background fit to the selected pixels.

Of course, sometimes a brightness gradient is desirable to focus attention on a foreground object and away from the background. This is easy to accomplish. Create a

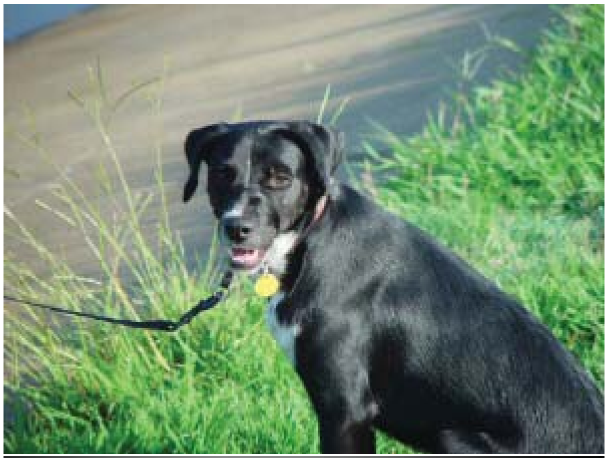

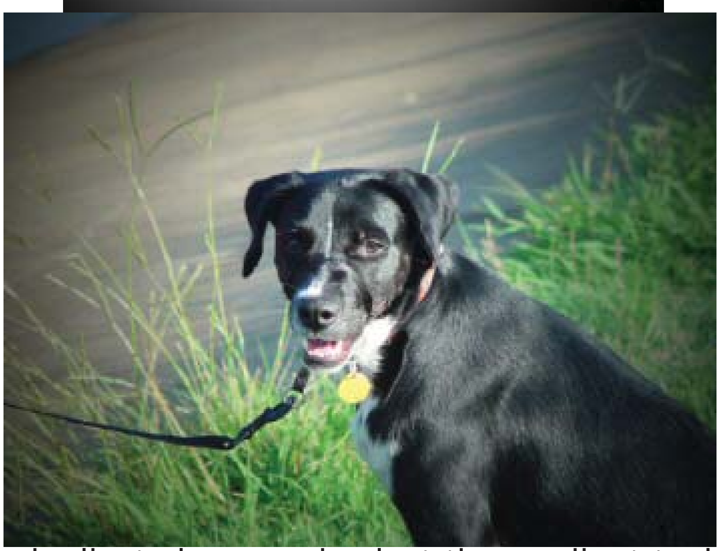

duplicate layer and select the gradient tool. Click on the center of interest (usually right between the eyes for a person or animal) for the white point of a radial gradient, and extend the gradient to the farthest corner of the image for the black point. Set the layer blending mode to multiply and adjust the opacity as desired to control the strength of the vignetting. Just as we use division to remove shading, we use multiplication to apply it. The illustration below shows an example.

Automatic shading correction: a) selecting the yellow shirts with the Select > Color Range function; b) removal of the background fit to the selected pixels.

In discussing human vision in the context of digital images, it's necessary to emphasize the importance of detecting edges and boundaries. This primary step is sometimes called the formation of a primary sketch from the image, highlighting the relationship between sketches and image understanding. The first few layers of neurons, right in the retina of the eye, use comparison between pixels and their neighborhood to find the edges, boundaries and detail that convey important information about the scene to higher levels of the visual cortex. Processing the stored digital image to help the eye find and extract this detail is one way that we can make images (or parts of them) more visually interpretable and more interesting. This is sometimes called "sharpening" the image, and sometimes is described as "edge enhancement."

The Unsharp Mask filter in Photoshop is the most widely used tool for accomplishing this purpose, but it is not the only one, nor always the most appropriate. By way of understanding what it does, and why it has such a counter-intuitive name, it may help to recall the original use of the method, in the photographic darkroom. Early film and print materials had a more limited dynamic range than most modern types, which made it difficult to capture the full range of contrast in prints, or the fine detail in both bright and dark regions of the image. This is still sometimes an issue since photographic prints (and computer hard copy) still have a narrower dynamic range and less tonal resolution than film or many digital camera images.

In order to use more of that range and resolution for the details in the image, at the sacrifice of some of the global contrast between bright and dark regions, film photographers could produce an “unsharp mask” by printing their image onto a piece of film material, at 1:1 magnification but slightly out of focus (hence “unsharp”). When developed, this film was a positive print, dense in regions that the negative was thin and vice versa. Sandwiching this “mask” with the negative and printing through both produced a final result in which the overall contrast between bright dark areas was greatly reduced, but in the regions around edges and detail the out-of-focus mask allowed the fine detail on the original negative to print through. The result was an image that used more of the print’s tonal range for the detail and less for the overall contrast.

It is easy to demonstrate this same process with a digital image, as shown below. A duplicate of the image is blurred using the Filter > Blur > Gaussian Blur function. The radius is adjusted to be just large enough to remove all of the detail that is important, typically a radius value of a few pixels. This blurred image still has the overall contrast between light and dark regions, but the fine detail and edges have been blurred. Subtracting this image from the original accomplishes the same thing as sandwiching the negative and unsharp mask together and printing through both of them. The result is the difference between the two images, which is just the fine detail and edges that were removed by the blurring. This detail is then added back to the original image to highlight those edges and that detail, which increases the visual sharpness of the resulting image.

The Unsharp Mask process: a) original image; b) blurred duplicate; c) image a minus image b; d) image a plus a percentage of image c.

The key variables in this process are the amount of the blur, and the amount of the detail image to add to the original. The Photoshop Unsharp Mask filter has sliders for both of those controls. The third control allows selection of a threshold value for how much difference there must be between a pixel and its neighbors to be considered detail and kept after the subtraction process.

Increasing the value will add sharpness to only the major contrast steps and not the places in the image where the detail has less contrast. Understanding what the sliders do in terms of the process will help the user to manipulate them to a successful result, but the criterion for selecting the “best” settings for any particular image is a visual one. If the detail that is of interest to the viewer is enhanced, then the image will have a better visual appearance and the operation is beneficial.

The Photoshop Unsharp Mask dialog, showing controls for the amount of the edge image to add to the original, the amount of blur to be applied, and a threshold for which edge detail to keep.

It is important not to make the blur in the unsharp mask too large. That increases the depth of the shadows or halos around edges, which increases their visual impact, but if the shadow from one edge or detail covers another, lesser edge it hides it and reduces the overall sense of detail and sharpness. The next two sets of illustrations show examples. In the first set, the spacing of the rock layers is fairly uniform and so a good setting of the unsharp mask improves all of the detail, while a larger setting degrades the image. In the next set, the detail occurs on many different size scales, and the best blur setting is the one that corresponds to the finest detail that is to be kept.

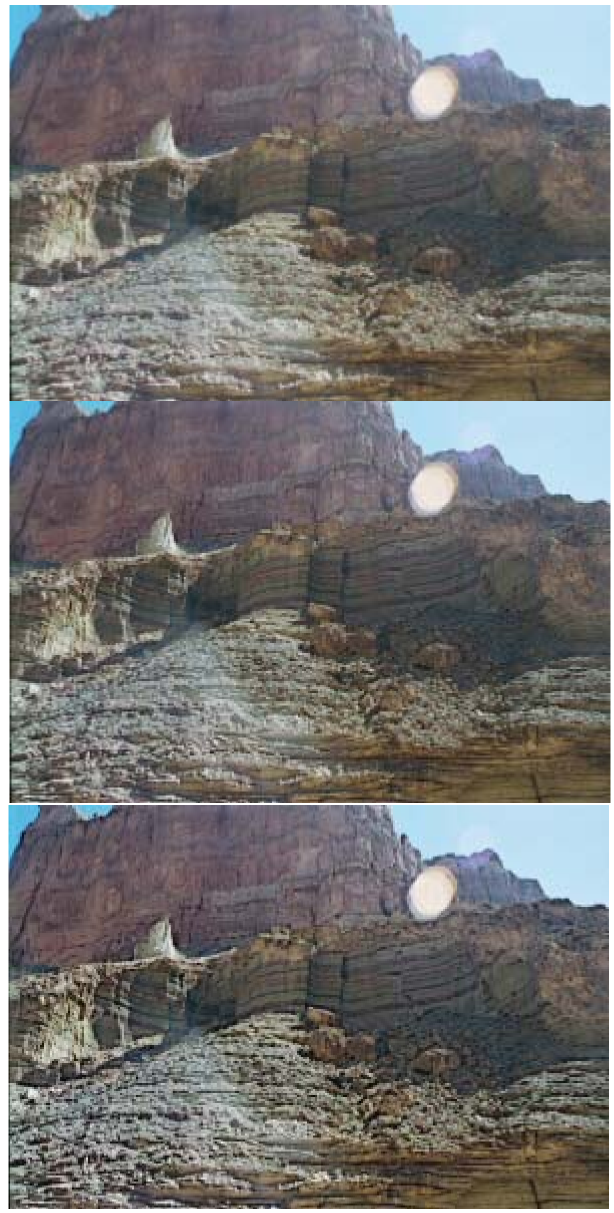

Varying the amount of blur: a) original image; b) good setting (2.5 pixels); c) too large a blur (9 pixels).



Varying the amount of blur: a) original image; b) good setting (2 pixels); c) too large a blur (10 pixels).

Applying an unsharp mask to a color image by processing each of the red, green and blue channels creates problems. The boost in edge contrast is a function of the pixel values on each side of the boundary, and except for neutral grey colors this is different in each color channel. The consequence is that the amount of change in each color is different, and this results in their combination producing a new color, or hue, that was not present in the image originally. Highlighting red areas with orange or purple, for example, creates a distraction for the viewer.

Noise in the image is also a problem. One common type of noise is simple random speckle superimposed on the image. This arises from statistical variations in collecting the electrons in the chip, amplifying them to produce a signal, and measuring the voltage produced, and is unavoidable. Lower light levels (requiring a higher effective ASA rating and more amplification) increase the amount of this type of noise. The unsharp mask responds to a difference between a pixel and its surroundings, but that is exactly what this type of noise is. Consequently, the visibility of random noise is increased by an unsharp mask, as shown below.

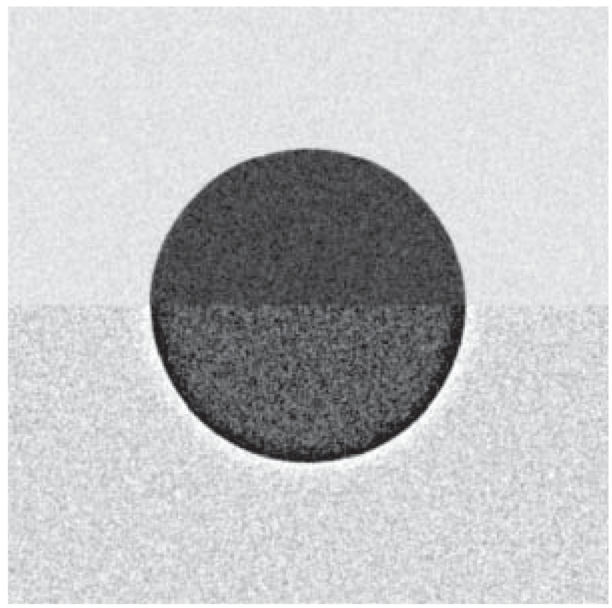

The visibility of the random speckle noise superimposed on this simple image of a circle is significantly increased by the application of an unsharp mask to the bottom half of the image.

But in a color image, the noise in each channel is largely independent of the others. There is usually the least noise in the green channel, because the most common sensor design, the Bayer pattern, uses more detectors for green and thus generates more signal. The blue channel typically has the most noise, because the silicon-based chips are not very sensitive to blue light (neither is human vision). Because the noise fluctuations in each color channel are different, at one pixel location the red value may be high and the green low, while at another the green may be high and the blue low. Each of these differences is amplified by the unsharp mask. When the variations in the different color channels are viewed together, they produce new colors, not present in the image originally. This colored “snow” is very distracting and degrades the image considerably.

As shown in the next illustrations, the unsharp mask works much better when applied not to the individual color channels, but to the intensity values only. This can be done in Photoshop by converting the image from RGB to L*a*b color space (Image > Mode > Lab), selecting just the Lightness channel, and performing the processing on that channel before converting the image back to RGB. This leaves the color values unchanged, so the random speckle noise in the various channels is still present but does not increase or give rise to new, distracting colors in the processing.

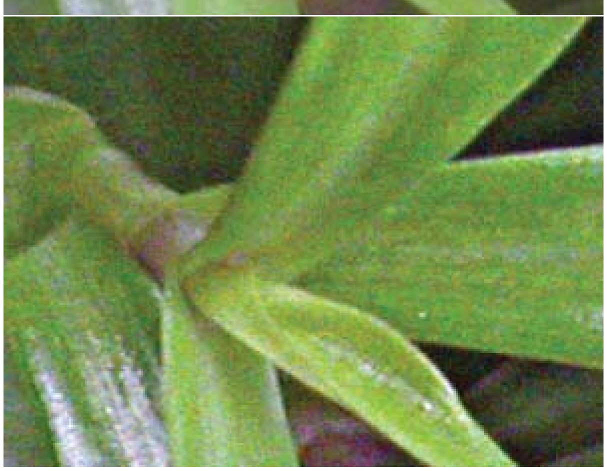

Sharpening a color image (enlarged to show pixel detail): a) original; b) unsharp mask applied in RGB channels; c) unsharp mask applied to intensity channel only. The visibility of the random speckle noise superimposed on this simple image of a circle is significantly increased by the application of an unsharp mask to the bottom half of the image.

Filtering the intensity channel only is a very general approach - almost all image processing operations should be applied to the intensity channel of a color image, not to the individual RGB channels. The exception is the reduction of noise, discussed below, which should be performed on each channel since that is where the noise arises and since it is largely independent in each channel.

Laplacian and Difference-of Gaussians Filters

The unsharp mask should be understood as one member of a family of operations that compare pixel values to find differences. On the very smallest scale, comparison of each pixel to its immediate neighbors (called a Laplacian operation in the image processing literature) is carried out with a filter as shown below. The Photoshop “custom” filter is a powerful tool for understanding how filters work and experimenting with various rules for combining a pixel with its neighbors.

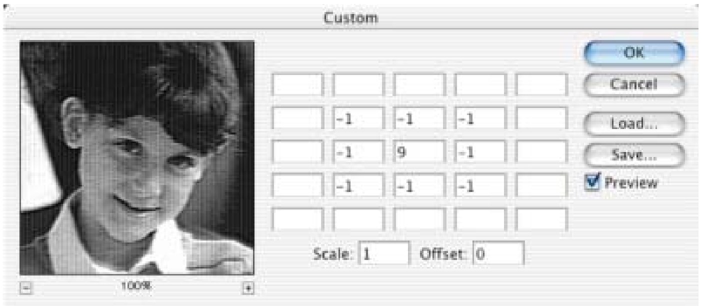

The Photoshop "custom" filter with weights entered that apply a classic Laplacian filter.

In the example, the process is to multiply each pixel value by 9 and then subtract each of the 8 neighbor pixels. In a region where all of the values are nearly the same, this leaves the original value nearly unchanged. But if there is a difference in brightness in any direction it produces a larger result with the effect that a pixel brighter than its neighbors becomes brighter still, and vice versa. That is the same effect produced by the unsharp mask, and it is easy to imagine the ring of "-1" values in the kernel of weights as the blurred mask that is subtracted from the sharp original. The Laplacian is a very good noise finder because it finds differences between immediate neighbors.

In contrast to the tiny comparison shown in the Laplacian, in the human eye there is a hierarchy of comparisons that find differences between a central region and its surroundings on many different scales. The filter that performs the same type of operation in the computer is a generalized version of the unsharp mask. Instead of subtracting a blurred copy from the original, we can make two blurred copies, one blurred just enough to smooth the pixel noise and the other enough to smooth out the important detail, and subtract one from the other. Because the blurring is accomplished with a Gaussian blur, this method is called the "Difference of Gaussians" (sometimes abbreviated DoG) filter.

The advantage of the DoG method when applied to a noisy image is shown below. By comparison to the first circle, the increase in edge contrast is not accompanied by an increase in visibility of the noise. The radius for the Gaussian blur used to reduce noise is typically in the range of 0.5 to 2 pixels, just enough to reduce the speckle, while the radius for the Gaussian blur used for the unsharp mask is typically several times greater. Of course, this assumes that the scale of the detail is larger than the scale of the noise.

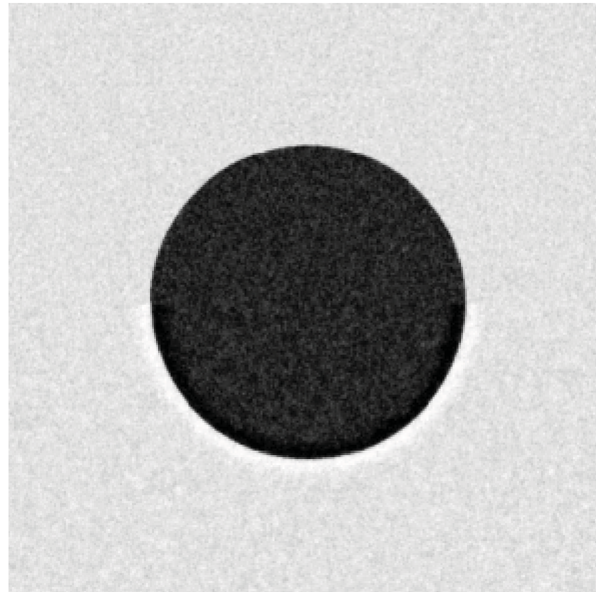

The random speckle noise superimposed on this simple image of a cirle is less visible than the circle drawn earlier.

The difference of Gaussians is a more general form of the unsharp mask, which includes two smoothing adjustments. One increases the edge contrast, as in the conventional routine, while the other reduces the noise. It would be equivalent to perform the smoothing first, separately, and then use a conventional unsharp mask filter except that it is difficult to know how much noise reduction is necessary until the sharpening process makes it more visually evident. The illustration below shows a plug-in (part of the Optipix plugin set from Reindeer Graphics) that implements this method. Setting the radius for noise reduction to zero produces the conventional unsharp mask result, which in this implementation is automatically applied to the intensity channel of a color image.

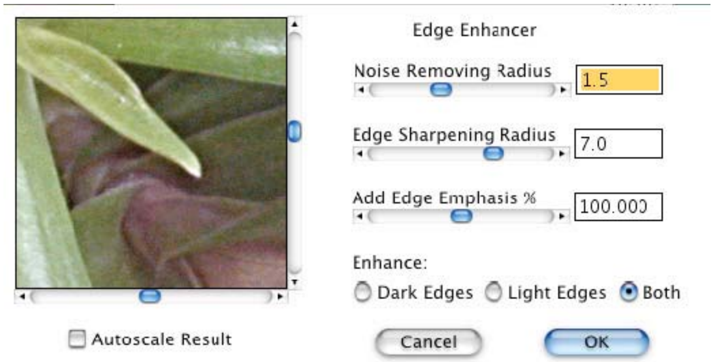

Edge sharpening using a difference-of-Gaussians method. The sliders allow interactive control over the radius of the Gaussians used to remove noise and enhance detail.

The other approach is to reduce the noise first. Since the noise in an RGB image is actually found in the individual color channels, and is mostly independent and has different characteristics in each channel, it is possible to apply a Gaussian smooth directly to each channel. In the example shown below, a different amount of smoothing was used in each channel (a radius of about 2 pixels in the noisier blue channel, a radius less than 1 pixel in the green channel, and an intermediate value in red). Then the conventional unsharp mask was applied to the intensity channel as described above. This produces the best result, but requires the most effort. And since the increase in noise due to the unsharp mask is hard to judge beforehand, a few trial-and-error iterations were needed to get the final result shown, reducing the noise without sacrificing detail sharpness.

Sharpening the earlier color image with noise reduction: a) difference-of-Gaussians enhancement (unsharp mask plus smoothing) applied to the intensity channel; b) smoothing each color channel separately, then applying an unsharp mask in the intensity channel.

The unsharp mask (and related methods like the difference of Gaussians) highlight edges and detail by increasing the change in brightness close to each step. The standard version of the technique, shown in examples in earlier columns, adds a bright halo along the bright edge of the step and a dark halo along the dark edge. Sometimes there are advantages in just using one or the other, rather than both.

For real images, this has the advantage of reducing interference between steps or detail and the haloes from other, nearby steps. Also, it can be used to make features stand out better from the background. To do this in Photoshop, duplicate the layer holding the image, apply the conventional unsharp mask, and then set the layer blending mode to “Darken” or “Lighten.” Of course, this only works for 8 bit per channel images (which can be put into layers), but the Optipix plug-in shown below, which allows direct selection of dark or light edges, also functions on 16 bit per channel pictures.

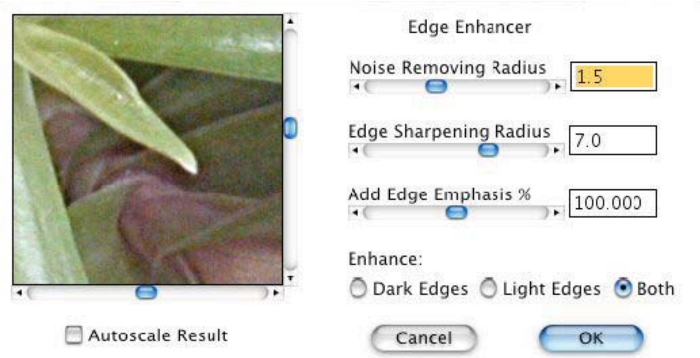

The general rules for deciding which edge halos should be used are: 1) light edges don’t show up well on light background tones, and vice versa; and 2) the halo should lie on the background rather than the foreground, to help the feature stand out without altering its brightness values. Sometimes these rules are in conflict and it is necessary to try several combinations (both edges, light edges only, dark edges only) to decide which is best. As well, different regions of the image may call for different answers.

In all of the examples shown to compare bright and dark edges, the width of the halos (radius of the blur function) and the amount of the effect added have been set larger than ideal, in order to make it easier to see and compare the effects. The two sets of images below illustrate the first rule. The light halos create visual artefacts (for instance next to the shutters) on the light colored house, and the dark halos merely interfere with other highlights on the dark dog. So we would use the dark edges on the light house and the light edges on the dark dog for best results.

Comparison: a) enlarged detail of original image; b) light edges only; c) dark edges only.

Comparison: a) enlarged detail of original image; b) light edges only; c) dark edges only.

The next two sets of images illustrate the second rule. The light halos are visually distracting and unnatural around the model’s eye, but the dark edges simply serve to increase the local visual contrast and bring out the existing detail. The leaf is darker than the background, so adding a light halo to the image helps the foreground object to stand out.

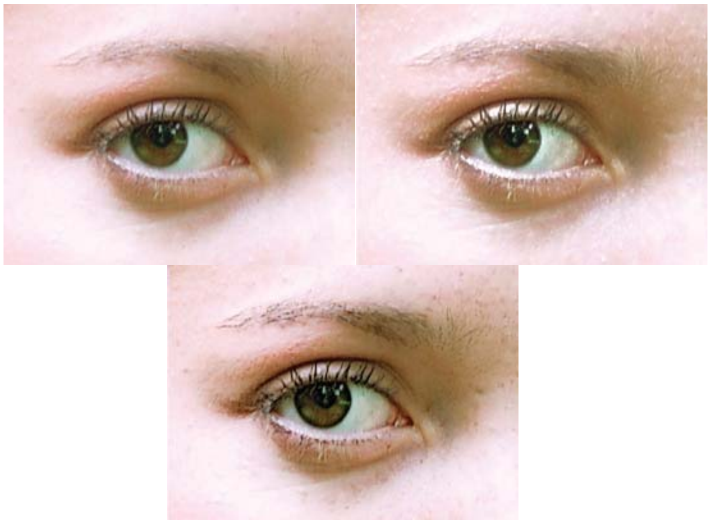

Comparison: a) enlarged detail of original image (using the image from the portrait further down); b) light edges only; c) dark edges only.

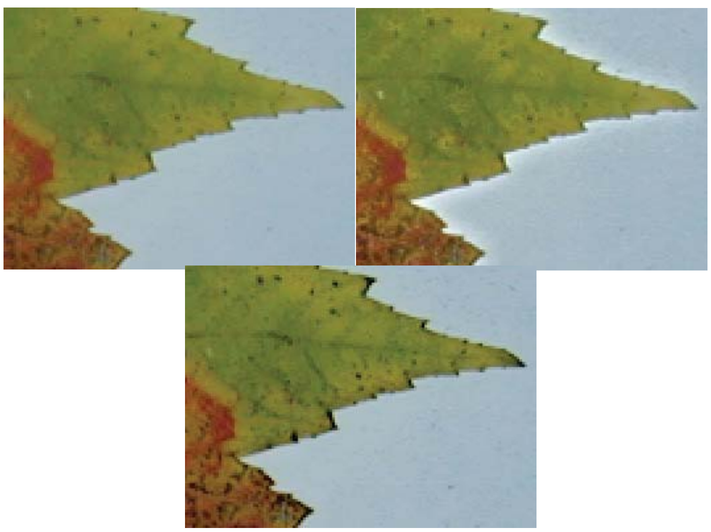

Comparison: a) enlarged detail of original image; b) light edges only; c) dark edges only.

The typical problem for users of the unsharp mask filter is adding too much additional local contrast. Many of the preceding examples have used settings for the radius of the Gaussian blur and the amount of the effect to add to the original that are large, in order to make the nature of the effect quite evident in the example images. Adding some local contrast can make a bland image into a good one. Adding too much can make a good one into a disaster. Seeing the effect on the computer screen and watching details “pop” out of the background can easily mislead the practitioner into the “some is good, more must be better” mindset that produces caricatures instead of photos. It is hard to dispassionately review the results of processing to detect this flaw, and it is ultimately up to the individual user to decide where the limit lies. And it should be noted that what looks best on the computer screen will not usually be the same as what looks best in print, because the printing process compresses contrast and blurs detail somewhat.

The next set of illustrations shows an example. The original picture has had some minor retouching of blemishes and darkening of the background, but is otherwise unmodified. A small amount of added local contrast using an unsharp mask (actually a difference of Gaussians) improves the vitality and sharpness of the photo. Increasing the radius and percentage settings by one third clearly blows out the highlights, creates artificial shadows, and hardens the image in an unfortunate way.

Taking the unsharp mask too far: a) original image; b) modest sharpening with good results; c) abuse of the technique.