Sharpening with Unsharp Mask in PHOTOSHOP

Introductory note and a personal observation.

This article was prompted by the remarks of a Judge at a recent competition, whose critique of many images correctly identified a dearth of post-processing skills, mainly in the area of sharpening.
I only became concerned when the judge recommended the use of a technique named 'High Pass Filter,' which is renowned for its aggression & is seldom used within the professional fraternity. Unless this filter is used with advanced image processing methods, such as masks & layer-blending, it can easily affect an image in an adverse way.
High Pass filter also needs to be applied in distinctly different ways according to bit size of the image and, the use of this filter by those unfamiliar with its aggressive nature, is a major cause of what is known as 'haloing' in images – the very thing that had been criticised as a major fault in many images.
There is no reason why you cannot use it, but the one I describe really is a much better and safer option.

Image sharpening in digital photography is not very easily understood by either photographers or those who 'post-process' images. So, what is sharpening? Why do we need it? What does it do? Why should you be applying it to your own images?

Unsharp Mask – What is it?

The unsharp mask is counterintuitively named, especially until you understand how it works, as explained later in this article. It was first used in Germany in 1936, well before the time of digital images. It was developed in connection with the printing of negatives through a darkroom enlarger.

Since that time, unsharp mask has become the professional industry's 'go to' tool. It is not only accurate but also very controllable. The algorithms used by it are also used (copied) by other sharpening tools, so understanding Unsharp Mask will give an insight into others.

So Why do we need sharpening?

Digital images, both those created in a digital camera, or by older technology such as scanning images to create digital copies of analogue work, will always suffer from limitations imposed on it by the size of the device's digital sensor.

Values of a modern digital camera sensor, measured against the human eye, performs very poorly. It shows just how limited current technology is - despite the fact that it has improved exponentially in the more recent years.

A human eye has been calculated to be the equivalent of a 580-megapixel camera sensor or CCD (i.e. Charged Coupled Device). It can distinguish lines of contrast with incredible clarity and sharpness. When compared to the current, so termed 'entry-level' camera, which has a CCD of around 20 megapixels.

Twenty megapixels is a whopping 30-fold ***decrease*** in quality, compared to the human eye. That decrease significantly restricts the number of data-points that the sensor can collect, resulting in our perception of it producing an overly soft or 'fuzzy' image.

The human eye sees in approximately 28 different shades of grey. This is significantly more than the 5 shades (*i.e. Black, quarter-tone, half or semi-tone, three-quarter tone and white*) that the old film cameras possessed. Presently, even our 'best' Digital Camera is only capable of distinguishing 13.5 shades of grey and most can only distinguish between 5 and 12 shades.

These limitations mean that when camera is facing a scene of a higher resolution than it is capable of capturing, *and it currently always is*, it is forced to capture the average of what it sees. It does this by using 'on-board' or 'in-built' software to severely compress the white and black 'ends' of the image. The result is the 'fuzziness' we associate with a lack of sharpness.

In effect, the camera has been asked to squeeze down a very large object into a very small box. Put simply, it is asked to squeeze an elephant into a shoebox. The camera is simply forced to do its best with the insufficient amount of data it is able to capture.

Demonstration

This digital mock-up should demonstrate the effect of our vision compared to that of a camera.

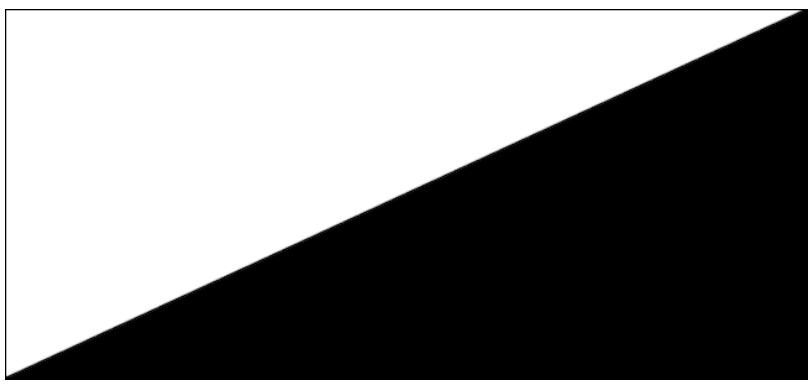

Figure 1 Triangles Viewed by the human eye

In figure 1, the space has been equally divided into two triangles, one black and one white. It is how the human eye would see it.

To us, it looks like just one crisp and continuous line between the black and white sections.

For the sake of demonstration let us pretend that the image in figure 1 is an intersection of two spaces in the real world. Perhaps two halves of a board that has been painted with such precision that, even with a magnifying glass, the line remains crisp and distinct. This is our maximum resolution and the human eye would perceive it as very sharp and crisp.

Now let us compare figure 1 with figure 2.

The line should be razor sharp, as in figure 1, but a lack of pixels causes problems for the camera's sensor. It cannot perceive it to be half white and half black

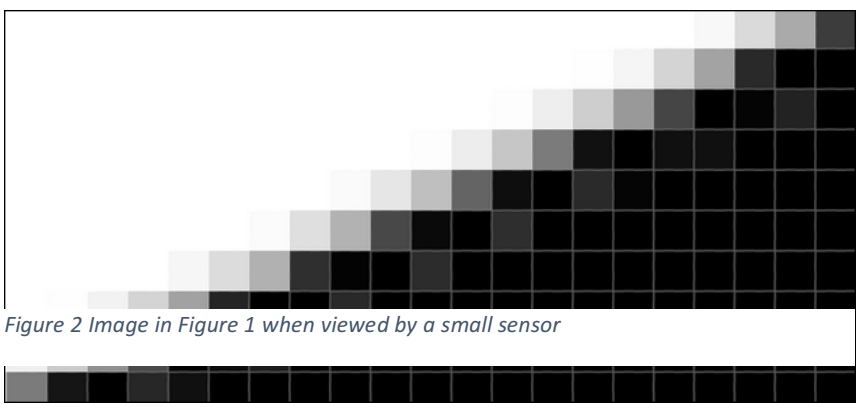

Figure 2 Image in Figure 1 when viewed by a small sensor

It cannot record a single value for the entire line of pixels. It is forced to use the average of the light, which is grey.

It cannot divide any single pixel into two colours, so it averages those at the division line to be grey.

The greater the pixel count a camera has, the more detail it can record but until this reaches the equivalent of a human eye, it will always come a point where the incoming data (that is, the light reflected from the subject being photographed) exceeds the capability of the sensor. The camera is thus forced to select a best estimate of shade and contrast between the edges, resulting in our seeing it as fuzzy.

(Note the step like appearance of the pixels. These are known as 'Jaggies.' See later.)

Using Unsharp Mask

The unsharp mask filter is almost counterintuitively named, until you understand how it works.

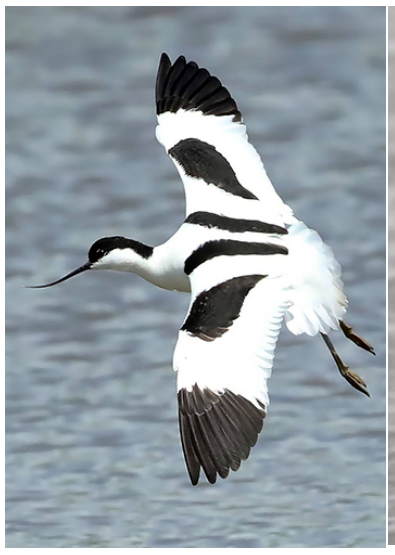

Figure 3 Avocet - original

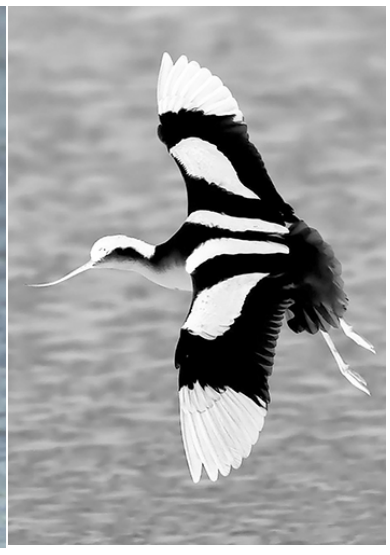

Figure 4 Avocet - inverted grey mask

Figure 5 Showing mask overlaid

From the original image, it creates a temporary mask which it then inverts (fig 4). This is placed over the original (fig 3) in overlay mode to locate the areas of high (sharp) and low (unsharp) contrast.

Figure 5 shows the resultant overlay. The areas shown in white are the ones to be 'sharpened.'

By using the fig 5 mask as a guide, the contrast in the unsharp areas is increased until the differences are equalized (50 percent value of the high low contrast pixels.) Thus, the unsharp mask tool is named as it indicates the unsharp areas - and corrects them.

Figure 6 Pet Dog

Using the dog in figure 6, we can demonstrate the effects of the unsharp mask filter and the adjustments we can make to it by using Adobe Photoshop.

N.B. The unsharp mask tool is also provided in other applications and is used in a similar way.

Note that the image at fig 6 is a small Jpeg that has not received any post-processing work save for that automatically done by the cameras in-built software.

There is nothing very wrong about the original image. The subject is centred and in focus.

Figure 7 Pet dog - close-up of original

If we zoom in, we look at the image in closer detail.

As we get really close in, the image looks softer than it should. This is a side effect of the way the image was processed in camera.

To sharpen the image and in preparation for using the unsharp mask filter, we must ensure it is displayed at either 100 percent or at 50 percent resolution on the screen.

This is necessary as anti-aliasing (*) algorithms have a tendency to distort the sharpening process unless these levels are used.

Note * ***antialiasing*** *is a software technique for diminishing what are called 'jaggies.' Jaggies are the stair-step lines that should be smooth (as seen in figure 2). They occur because the device – the camera, monitor or printer, has insufficient resolution to represent a smooth line.*

In Photoshop, goto - Filters ->Sharpen -> Unsharp Mask to get this.

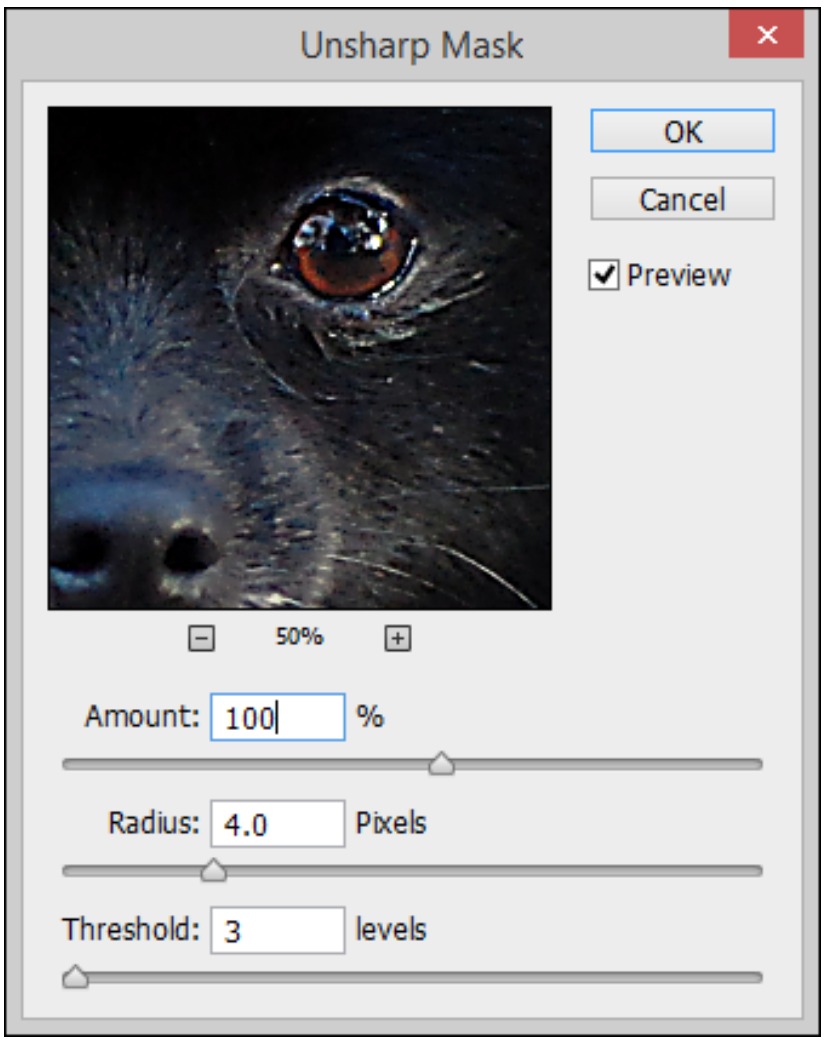

Figure 8 Unsharp mask controls

The unsharp mask tool (see figure 8) is fairly universally used and there are three settings by which it can be altered - Amount, Radius, & Threshold.

The easiest way to understand these settings is to play around with them. The key points are listed below.

1. Amount: Is always listed as a percentage. It indicates the degree of adjustment (how much lighter the lighter edges are to get, and how darker the darker edges are to get). The range is extreme, so somewhere between 50-100 percent is a safe place to start for an 8-bit image. This may be safely increased significantly for a 16-bit image and it is not unusual for this to be over 200%.
2. Radius: Specifies how large the area around each corrected point the effect is affected. The radius and the amount are intertwined; if you reduce the Amount you can increase the Radius value (and vice versa). Increasing either or both to high values will result in distortion to the colour & contrast (questionable as an artistic effect but it certainly will not make for a very natural looking image).
3. Threshold: The threshold function determines where the sharpening algorithm will be applied based on a minimum brightness/contrast level. It is best set to zero for general work. (*I personally only ever use it set to zero. I completely ignore this control.*) It can, occasionally, be useful for selectively increasing the contrast in high contrast areas, especially around the eyes, without over-sharpening smooth areas such as facial skin. The lower the value the more the image will be uniformly sharpened. Remember – it is inverted so 0 is 100% & vice versa.

Figure 9 is a representation of a typical sharpening curve.

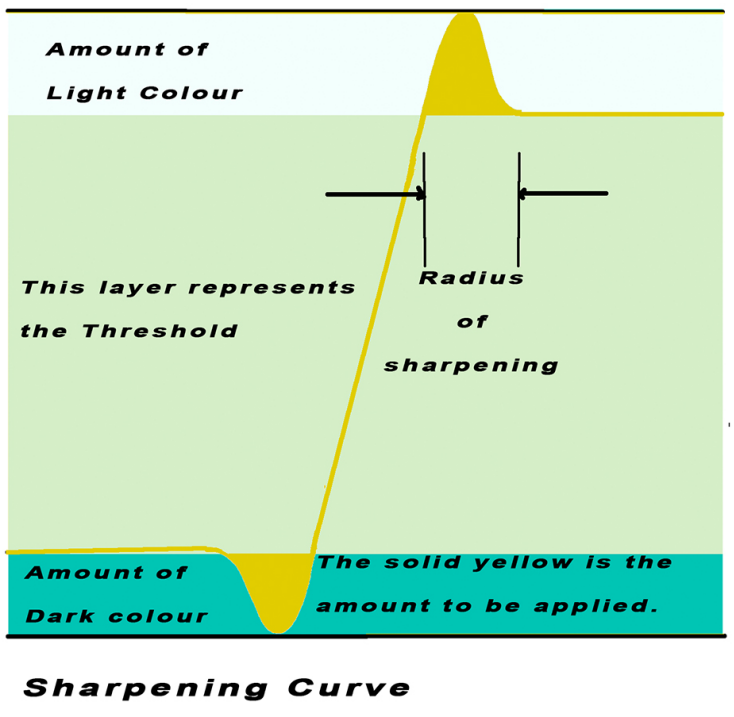

Figure 9 Sharpening curve

In Figure 9, the top is where the lighter pixels are. They are sharpened by increasing their luminosity or brightness value. The bottom is where the darker pixels are to have the opposite effect applied to them.

Prior to sharpening, in the actual image, the narrower yellow line represents the division of the lighter and darker areas. The solid yellow part is where the width or radius and its depth, are to be sharpened and the amount by which this is to be done.

Applying the settings, we have shown in figure 8 (i.e. 100/4/3) will significantly sharpen the image shown at figure 6, as shown in figures 10 and 11.

Figure 10 Post processing of figure 6

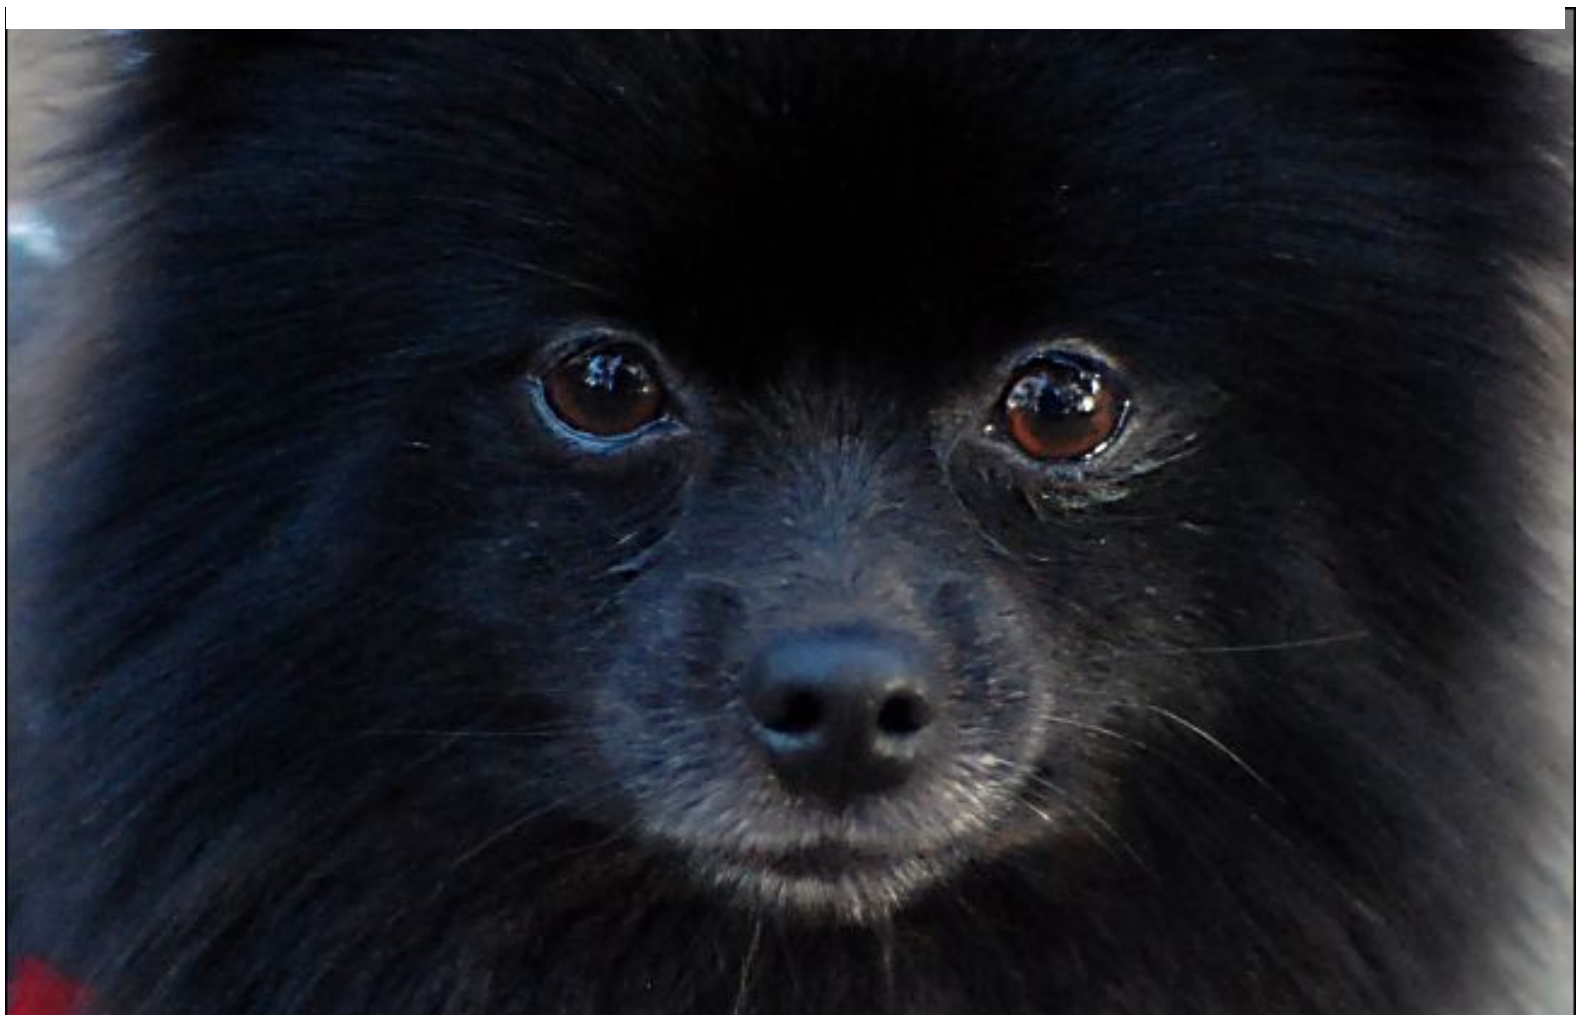

Figure 11 post processing - close -up

The changes from the image shown as figure 6 to those in figures 10 & 11 are self-evident. The eyes are sharper and the muzzle is more distinct.

The goals of sharpening an image.

The objective of sharpening any image is to recreate a sharpness that a human eye would have seen. It should never be used to create an intense and noticeable contrast that identifies the fact that image correction (or post-processing) has occurred – leaving the viewer question what kind of manipulation has been done.

What sharpening cannot do to an image.

While looking at the images used in this article (figs 6.7.10 &11), it's important to stress what the unsharp mask tool ***cannot*** do as well as what it can.

It does make the area that is in focus within the image look much better by sharpening up the edges and giving it a crisp look. What it cannot do is add any detail that was not recorded by the CCD, does not exist, or is too out of focus in the original image.

In figure 9 we have shown how the eyes, muzzle, and nose have become sharper, as did the facial fur, but the lead, concrete, moss, and leaves were all unaffected by the process. Those objects were so far out of focus that no amount of sharpening could have ever created even an illusion of them being sharp.

What Actually happens to an image when it is sharpened

The image at figure 12 is one which has been expanded to show the actual pixels involved in the sharpening process.

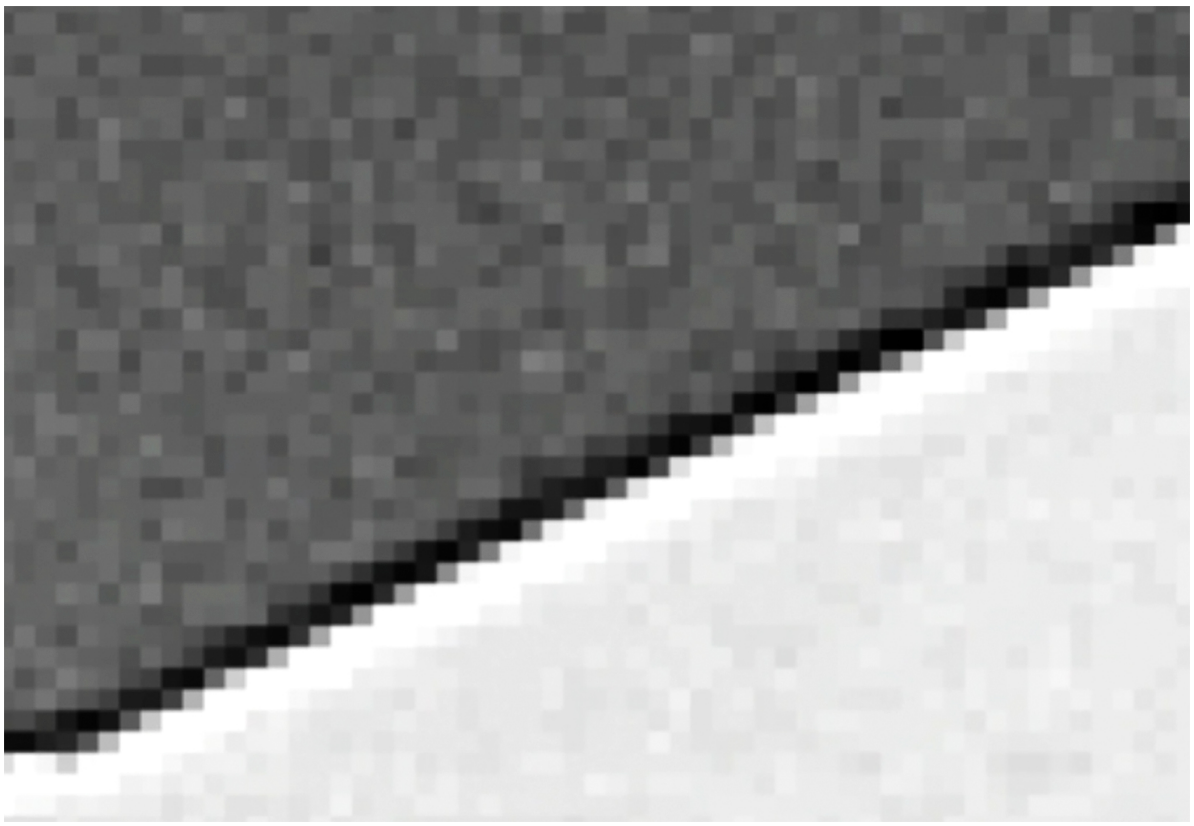

The image has been inverted and thus the darker triangle represents the background which, in the real image was the sky. The lighter triangle was that formed by a part of a bride's arm.

Figure 12 Result of Sharpening

Note how the 'sharpening' has seemingly 'drawn' a black line two pixels wide on the 'sky' side of the image (remember, it is inverted.) A white line of equal width is also on the edge of the arm. Both lines have been blended (feathered) slightly towards their opposite colours.

One could easily be tricked into believing that both the white and black lines have been painted between the colours that separate the light and dark areas – the white on the dark side of the delineation and the black on the other.

By using advanced processing techniques, which is beyond the scope of this article, it is possible to either increase or decrease either the amount of black or white pixels shown in figure 12. Accordingly, should the image warrant it, the image could for example, contain a black line that is three pixels wide and a white line of only one pixel wide. This is especially useful when photographing such items as jewellery containing diamonds and the like as it would be detrimental to show a brilliant stone which is surrounded by a heavy black line.

Unsharp Mask Tips and Tricks

Although people and animals benefit well from application of the unsharp mask it can also help to make any image stand out more.

Figure 13 Coffee Beans

Comparison of the two halves of the coffee bean images at figure 13 shows that there is little wrong with the original on the left, save for a little 'softness, once it has been corrected using unsharp mask, it gives on a more pleasing and realistic appearance as shown on the right. It becomes crisper, without becoming too 'over-blown.'

Unsharp Mask Tips and Tricks – consider the following: -

To get the most out of unsharp mask the following may help to ensure that the sharpening process is a smooth one.

Disable in-camera sharpening.

Above all else you want to disable in-camera sharpening. Point and shoot cameras almost always have on-board sharpening whereas high-end DSLR cameras rarely do. It seems that camera manufacturers presume that the point and shoot user will not be doing any post-processing work, but the DSLR almost certainly will.

Do Not Sharpen JPEG images

Sharpening works best with RAW and those files containing large amounts of data, often referred to as 'Lossies.' These include common formats such as TIFF; PSD; PSB & DNG etc.

Inexperienced and novice photographers seldom realise that images shot in 'JPEG' format have already been heavily processed within the camera, which includes sharpening. Most of the original information recorded, from which the JPEG had been made, has been 'un-recoverably' deleted by the camera and is unavailable for further processing.

Consider JPEG images to be the same as those that referred to above. Any image that has, in effect, been ‘double-processed’ with any form of sharpening will tend to look pretty awful and unrealistic.

JPEGS may be sharpened slightly, but extreme care should be exercised when doing so. They have, by definition, already been compressed by the camera’s software by a factor of ten compared to that of a RAW file.

Focus is king.

A crisp physical focus in the camera is worth ten times more than any sharpening tool can ever give you. Perfecting focus skills and using only lenses which fit well and are not ‘soft’ in focus will help make images sharp.

There is no magical way in which the unsharp mask tool can expand or fix the focal plane of an image. Remember - you can only sharpen what is already in focus.

Less is more.

Use the unsharp mask ‘just enough’ to give the image a little lift.

Think of it as the comparison of a High-definition (1080p) television with one that 4K resolution. 1080p is good when compared to previous televisions, but the 4K has a sharpness that just bursts right out of the screen.

When adjusting or comparing images, the aim is get ‘just-right’ sharpness. One that takes you from you thinking “That’s nice” to the one where you think “Wow, that really is crisp.”

There is a fine line between the two, so do not over-do it. Increasing the sharpening too much always leads to an unnatural looking image.

Sharpen last.

Although many experts in post-processing often sharpen their work in a few short stages, most leave their main sharpening until after all other editing work has finished. Colour adjustment, fixing dust, or stuck pixels, or anything else should be done first. You can always achieve a better result by leaving the main sharpening process until last.

Remember that images displayed by a projector will not need as much sharpening as one that will be printed. Projected Digital Images (PDi’s) have, by definition, light ‘being forced’ through them. That this alone will make it easier for the human eye to differentiate between the light and dark areas in the image.

Think of sharpening an image as polishing a piece of jewellery after you’ve finished working on it. It’s the last step after every piece is placed, every bit of metal bent and soldered, and the it’s ready for the gallery.

Check your work

Making what is known as a ‘stamp layer,’ before the final sharpening will allow you to see exactly what has been sharpened and how good it is.

Make a ‘stamp layer’ by using the keyboard shortcut Command/Shift/Option E on a mac or Control/Shift/Alt E on a Windows based computer. Toggling the layers ‘eyeball’ on and off after sharpening will show which pixels have been effected, especially if you have ‘zoomed’ up the image size on the screen. Try to get it to the size shown in figure 12 – but realize that this isn’t what you are going to see at the ‘normal’ resolution level (Command/Control zero) – you should hardly notice it there.

Remember, if the sharpening does seem harsh that you can affect it by reducing the opacity of the layer or, by using a different layer blending mode such as Luminosity, or even both. If you are brave enough, you could even try blending the final ‘stamp’ layer with those beneath it – as any experienced post-processor would. Just double click on the layer and use the sliders on the blend tool to get what you want.

Conclusion

The Unsharp Mask may be a very old tool in photographic terms, but I would wholly recommend its use. It really hasn’t ever been bettered. That is not only my personal opinion, but that of many post-processing experts also. It is why it is so widely used by most post processing labs around the world. Why not give it a try – what is there to lose?

The foregoing has been based on my personal experiences, gained during recent training for the Adobe Certified Expert (ACE) qualification, and on an article published by Jason Fitzpatrick. Some information was also provided by Wikipedia.

Links to articles: -

https://www.howtogeek.com/215920/htg-explains-what-digital-image-sharpening-is-and-why-you-should-be-doing-it/

https://en.wikipedia.org/wiki/Unsharp_masking

Terry Wagg