
Channel Surfing
A Guide to Image Channels in Photo Editing Software
© Seth D. Madell, November 2019

About this Guide…

This guide is a compilation of a 5-part series that was originally published in the July through November, 2019 issues of *Luminosity* - the Doña Ana Photography Club newsletter.

All of the chapters were written by Seth D. Madell, who served as newsletter editor throughout that period.

Channel Surfing – Chapter One

Your journey of understanding begins with Color Channels

Your photo editor probably has a panel called "Channels." You've probably seen references to channels in multiple places when you're creating and editing photos. Maybe you've skipped right over those panels and sliders, believing that channels were one of those esoteric things only geeks and nerds needed to understand. Turns out, understanding channels can help make your photos even better than they already are.

This is Chapter 1 of a five-chapter series about image channels. I have split this into five separate chapters because the entire thing, all at once, would simply be too long. Also, for those who are complete channel newbies, this is a lot to digest.

In Chapter 1, we will explore color channels. We'll talk about how your pixels are actually represented by combinations of different colors, and how color channels are a representation of those combinations. In Chapters 2 and 3, we'll introduce alpha channels. Alpha channels are infinitely more geeky than color channels, but are the other piece of the puzzle - just as necessary and even more powerful. In Chapters 4 and 5, we'll try to introduce some real world uses for all this stuff. (Honest, it isn't just a mental exercise.)

It's also necessary to state up front that, for simplicity's sake, I'm going to assume that all of our images are 8 bit images. Although there are plenty of reasons to edit in 16 bits, or even in 32 bits, explaining channels will be easier if we stick to the simplest type of color images. In truth, everything we talk about can be scaled up to explain greater bit depth photos; it's simply not necessary for these articles.

Introduction to Color Models

Photo editors are also known as "raster editors" because they display, and let us interact with, images that are made up of individual pixels. In the process of manipulating our photos (even if it's just resizing or cropping) we may be changing those pixels. Each of the pixels in our image has a color, and our software needs a way to show us what those colors are.

Colors are described using "color models." Although there are quite a few of them, the two that are used the most are the RGB model, and the CMYK model. The RGB model uses combinations of Red, Green, and Blue to represent colors. The CMYK model also uses combinations to represent colors, but uses Cyan, Magenta, Yellow and Black. The RGB method is additive, meaning that different amounts of red, green, and blue are added to each other to create the desired color. The CMYK model is subtractive, meaning that cyan, magenta, yellow, and black subtract from each other to create colors.

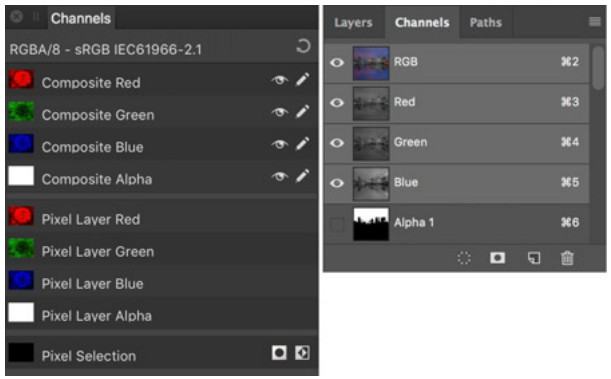

The Channels panel in Affinity Photo (left) and Photoshop (right)

Additive color models are used when describing colors in a system based on light. This is handy when we're putting a color up on a monitor. Three different colors of light can be added to create the colors in our photos.

Subtractive color models are used when describing colors in a system based on ink. These methods are used when creating colors on a printer.

Channels are like Spreadsheets

Color channels are the way that our image editors store information about the color of each of the pixels in an image. Let's limit our talk to RGB images. If we assume that every single one of the pixels in an image is some combination of red, green, and blue, then we really only need to store 3 numbers to describe the color of a pixel. Each pixel in our image has an amount of red, an amount of green, and an amount of blue. In each case, that amount is a number from 0 to 255 (since an 8-bit image has 2^8 possible shades of each color).

Think of the image as a matrix of pixels – each pixel has a position in the image, x number of pixels from the left, and y number of pixels from the top. If we created a spreadsheet that was x columns wide, and y rows tall, then we would have a grid of cells that could contain a number for each of the pixels in the image. If we did this 3 times, we could enter the values of red, green, and blue into each of the spreadsheets.

Now imagine, if you will, that our Red pixel spreadsheet (which contains values between 0 and 255 in each of the cells corresponding to the image's pixels) was itself made into an image. Since there's only 1 number in each of the cells, the image could only be a monotone image. In other words, a single "spreadsheet" of pixel values could only display a single color in 256 different shades. Typically, this is the way we form greyscale images - shades of grey varying from black (those with a pixel value of 0) to white (those with a pixel value of 255).

So, each of the colors that make up an image (red, green, and blue) can themselves be represented as a greyscale image. And that is what you see when you look at a color channel. The channel is represented by the software as a greyscale image, with each pixel in the channel displayed as a shade of black and white based on how much of that color is making up the various pixels of the total image.

By convention, a value of 0 is represented as black; a value of 255 is pure white.

Show Me an Example!

Consider a picture of a Red circle. To start with, we'll put the red circle on a Black background. All of the pixels in the circle have a red value of 255. That is, the circle is pure red. However, the pixels in the circle have green and blue values of 0, since there is no green and no blue in those red pixels. Also, since the circle is on a black background, the rest of the image has red, green, and blue

values of 0 (since there is no color whatsoever in those areas). Seen in the right margin is the image itself, and also the greyscale representations of the Red, Green, and Blue values.

Let's expand that example to include an image with 3 circles - one pure red, one pure green, and another pure blue. We'll show the original image along with the red, green, and blue channel images below.

Now let's change the example to include a colored circle that is a combination of red, green, and blue. We'll draw a circle whose color has a Red value of 210, a Green value of 130, and a Blue value of 80. Also, let's now put this on a White background; white is a combination of pure red, green, and blue, so it has a value of 255 in all of the color channels.

The circle's color has a lot of red (a value of 210 out of 255), so it is represented as a very light shade of grey (which is closer to white) in the Red Channel. The circle's color has an intermediate amount of green, and a low amount of blue, so it has a mid-grey representation in the Green Channel image, and a very dark grey representation in the Blue Channel image. Also note that all 3 channels are surrounded by white, since the original image has a white background, and white is made up of pure red, green, and blue (all with values of 255).

Take Me Back to the Real World!

Let's now look at a real image, made up of many different colors. Shown on the right is the original photograph, along with the greyscale representations of the Red, Green, and Blue Channels. Notice how the red areas of the photo are mostly white in the Red Channel; the green and blue areas of the photo are mostly white in the Green and Blue Channels, respectively.

Put this Together for Me

Don't think of channels as being the same thing as layers. They're not. Color channels are like taking a single layer and creating 3 separate images out of it – one in each of the primary colors. Color channels let us visualize the components of our image's pixel colors. They also hold some of the keys to manipulating those colors when editing.

If you thought this was fun, wait until Chapter 2 introduces Alpha channels. Not only can we talk about color, we can also talk about opacity, selections, masks – all the stuff that "real geeks" need and want to know!

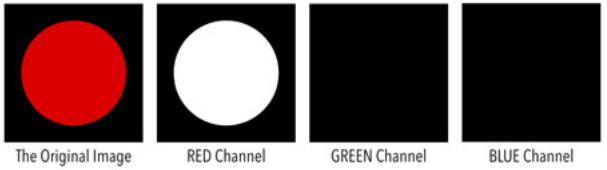

RGB Channels of a Circle on a Black Background

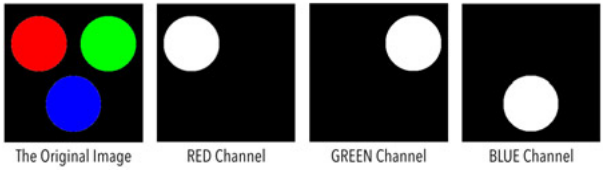

RGB Channels of Multiple Primary Color Circles

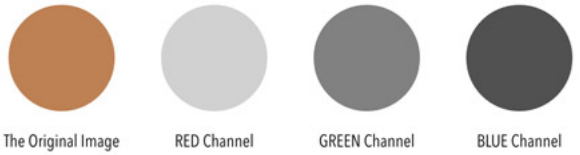

RGB Channels of a Non-Primary Color Circle

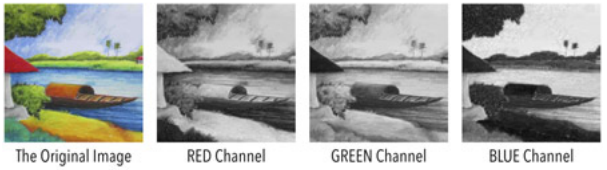

RGB Channels of a Real World Image

Channel Surfing – Chapter Two

To understand alpha channels, first understand opacity and selections

In Chapter 1, we introduced color channels. In RGB images, each of the pixels is made up of red, green and blue in specific amounts. We saw that we can represent each of the primary colors as a greyscale image, and that each layer can be broken up into 3 of these monotone images – one for each of the three colors.

Remember that our Red circle on a black background had RGB channels that looked like this (right). The white area was the circle made up of pure red, while the black area is the background which contains no red at all. In this case, the channel's monotone image is a sort of on/off switch for the primary color Red. Pixels that are pure red are active - the "on" state; pixels with no red are inactive - the "off" state. The Blue and Green channels are completely black, since they contain no red at all

The truth is that pixels can be more than just on or off – active or inactive. They can be partially active, too. Remember our last example with the circle that was made up of Red=210, Green=130, and Blue=80. The RGB channels had greyscale circles that were all shades of grey – none were completely white or black. For each of the primary colors, the circle's pixels were partially on and partially off.

Introducing Opacity

You probably know that a layer has a setting called "Opacity." When you set this to something other than the default value of 100%, the layer starts to become more and more transparent. You can start to see through the layer, where the image(s) underneath it become more and more visible.

Imagine we have a document that contains a Purple layer on the bottom, and a photo of a man in a room on top. Normally, since the photo of the man has (by default) 100% opacity, we see the top layer but not the purple on the bottom.

If we lower the opacity of the top layer, however, the purple layer starts to peek through. What we've done is made the photo of the man slightly transparent, and this lets us see through the photo in order to look at the purple layer underneath.

Let's lower the opacity of the top layer (the one with the photo) to 60%. When we do this, we can still see the man in the room, but the purple layer underneath him starts to become more obvious. We're partially seeing through the top layer, and viewing the bottom layer.

What changed? Remember, at the very beginning of last month's article, we looked at the Channels panel in Photoshop and Affinity Photo. Maybe you remember that there were Red, Green, and Blue channels. But did you notice the extra channel called Alpha? What is that, anyway?

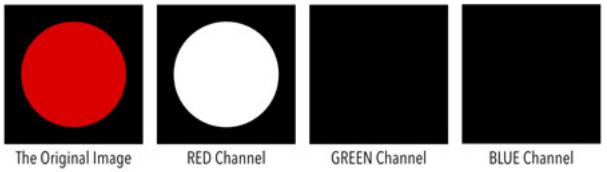

RGB Channels of a Red Circle on a Black Background

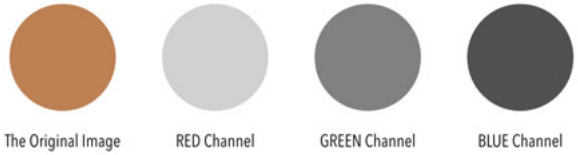

RGB Channels of a Non-Primary Color Circle

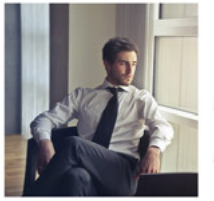
Photo - Man in Room

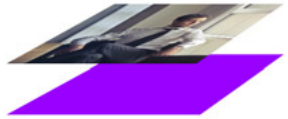
Photo layer above Purple layer


Original Layer Stack

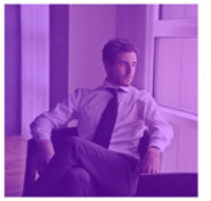
Opacity of Man at 60%

Alpha channels are traditionally said to contain information about transparency. I think the better way to understand them is to say that they contain information about Opacity. Opacity is really just the opposite of transparency – a layer that has an opacity of 60% has a transparency of 40%. But, in the language of channels, a pixel that is "on" or "active" is one in which the channel is acting on that pixel. A channel that describes Opacity (rather than Transparency) implies that when that channel is active, the pixel is opaque (not transparent), and this is how Alpha channels work.

Put another way, the pixels in our RGB image have four channels to describe them, not three. The Red, Green, and Blue channels combine to describe the color of each pixel. Adding information from the Alpha channel describes how visible that pixel is.

A pixel can have any given color (described by its color channels – red, green, and blue) and can have any degree of opacity/transparency (described by its alpha channel). The combination of the color channels and the alpha channels create a complete description of each individual pixel.

When a layer, like the photo of the man, has 100% Opacity, it has no transparency. Its alpha channel is completely white, since its opacity is 100% active. When a layer has 0% opacity, it is completely transparent – invisible – and its Alpha Channel is completely black, since its opacity is inactive.

When a layer has an opacity of 60%, its Alpha Channel is 60% white and 40% black (or, just lighter than middle grey) because the opacity setting of the pixels is only 60% active. Here is the Alpha Channel for the top layer of our photo stack (the photo of the Man in the Room), which has a 60% opacity.

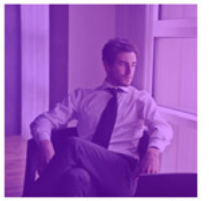
Opacity of Man at 60%


Alpha channel Man layer

Changing Opacity in Areas of a Photo

Now, imagine that you started with the photo of the Man in the Room. Leave the opacity set to 100%. Forget about the purple layer, for the time being. But, take the Eraser tool and erase some of the pixels on the left side.

If we inspect the Alpha Channel for this image, we'll see two areas distinct from each other. The majority of the channel is white, since this represents the portions of the photo that remain visible – everything that wasn't erased. However, now there is a squiggle down the side that is black. This black area represents the portion of the photo we erased, where the pixels are no longer visible. Because we erased all color information from those pixels, they have taken on 0% opacity (complete transparency).


Man in Room Photo


Erase a Portion

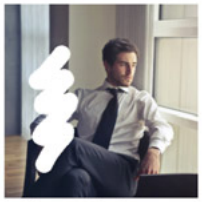
Photo with Erasure


Alpha Channel

Selections are like Temporary Alpha Channels

Before we go any further, let's take a moment to talk about Selections. You know that you can use lots of different tools to make selections. These may include the lassos (freehand, polygon, magnetic), various marquees, the selection brushes, and so forth. When you make a selection, you

have segregated your photo into distinct areas – those pixels that are within the selection and those that are outside of the selection.

A selection is a kind of Alpha channel. It is a special case, in that the selection is a temporary state. You can de-select all or some of a selection without actually changing anything in your photo. There is nothing permanent about a selection.

A selection is simply a temporary on/off switch for areas of your photo. (And, yes, selections can be partially on and partially off, just like alpha channels. We sometimes talk about pixels which are "partially selected.")

Sometimes, your software may have a feature called a Quick Mask. This is a way of visualizing selections, similar to the way channels let us visualize primary colors or opacity. When we make a selection and invoke the Quick Mask, the software typically shows us the selected area as one might normally see it, and the un-selected area with a red overlay (similar to the Black area in an Alpha channel representation). The normal area (the area without the red overlay) is the area of the active selection; the other parts of the photo (those covered by the red overlay) are inactive – that is, not selected.

Original Photo of Man Quick Mask of Man

Many software packages will allow us to Save a selection, so that we can use it later. Maybe you've spent a long time creating a selection of the subject of your photo (a person, a sky – anything, really) and you'd like to be able to call back your selection later on. When you save the selection, you can typically recall it as a selection by choosing it from a list; your selection reappears just as it was when you first created it.

Surprise, surprise – saving a selection means creating a new Alpha channel. This particular alpha channel does not affect the pixels of your photo, unless you use them to do so later. It is added data, above and beyond the information in your photo when you started.

Here's our photo with just the man selected. And, we've saved that selection as a Spare Channel in our Channels panel. You can look at this channel just as you would any other channel, and the greyscale representation shows you the Active portion (the man) in white and the Inactive portion (the rest of the photo) in black.

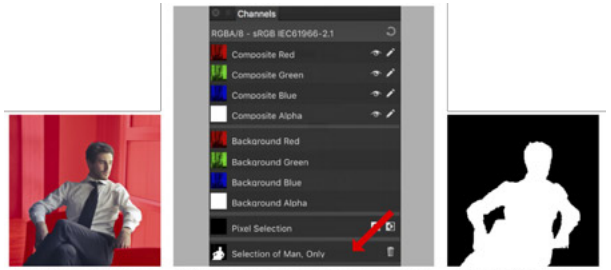

Quick Mask of Man Selection, saved as an Alpha Channel Spare Alpha Channel

Where is All of This Going?

In Chapter 3, we'll talk about Masks. Masks are another special type of Alpha Channel; they are the way that selection, opacity, and other aspects of alpha channels all come together.

Masks can reveal and conceal different parts of a photo by changing the opacity of the pixels in the areas you choose. Masks can help apply adjustments to some areas of a photo, but not others, by varying the opacity of the adjustment. Masks let us stack multiple photos on top of each other, allowing different parts of each photo to be visible when the entire layer stack is viewed.

Masks are the difference between good editing and great editing. Are you having fun yet?

Channel Surfing – Chapter Three

The one about masks, where black conceals and white reveals

When we started this series, we began by introducing color channels - the three separate parts of an RGB image that came together to specify each pixel's color. In Chapter 2, we introduced alpha channels – a fourth parameter that can specify the opacity of each pixel. We started by talking about the opacity of a layer, and then talked briefly about selections as a kind of temporary alpha channel.

In this chapter, we're talking about Masks. Masks most often reside in a separate layer and represent an additional alpha channel. When added to a layer, they make portions of that layer transparent.

Since a mask can be added as a separate layer, it has all of the non-destructive benefits that extra layers have – they can be edited separately, turned on or off, or deleted without changing the layer they affect.

How Does Software Represent a Mask?

Photo editing software will usually show you an additional layer, clipped ("attached," as it were) to the photo or other layer it is affecting. Photoshop shows you a separate icon next to the thumbnail in the Layers panel. Affinity Photo shows you a separate mask layer as a "child" of the photo layer. Other software packages may represent this differently.

It is important to remember that, in most circumstances, a mask affects only the layer to which it is attached. It does not affect other layers.

What Does a Mask Do?

An RGB photo, placed into your Layers stack, has 4 channels. Three of those channels (red, green, and blue) determine the color of the photo's pixels. The fourth channel (alpha) determines the opacity of each of those pixels. Most photos, especially those that have not yet been edited (such as files taken from a digital camera), have alpha channels which are pure white. That is, there are no pixels in the photo which are anything other than 100% opaque.

When we want to change the opacity of some, but not all, of the pixels in our photo, it is easiest to do that in a separate layer. This layer, called a Mask, represents an additional alpha channel. Since it is attached to a specific layer, and affects only that layer, it is an easy way to change the opacity of that layer's pixels without editing the layer directly. This is "non-destructive editing," since the photo itself remains unchanged.

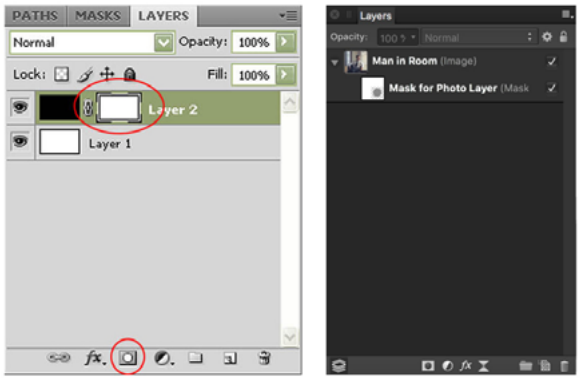

Mask Placed on Photo Layer
in Photoshop in Affinity Photo

The Simplest Masks

The simplest masks are those which are only black or white. These masks have well-demarcated areas, without any "fading in" or "fading out." Since masks are alpha channels, remember that the white areas are 100% active (and, therefore, completely opaque); the black areas are 0% active (and are, therefore, completely transparent).

If I attach a mask to a photo, and draw a black circle in the middle of the mask, I am effectively telling the photo that the pixels corresponding to the black circle in the mask should be rendered at 0% opacity. The black areas of the mask have caused a change in the visibility of the photo layer without changing the photo layer directly.

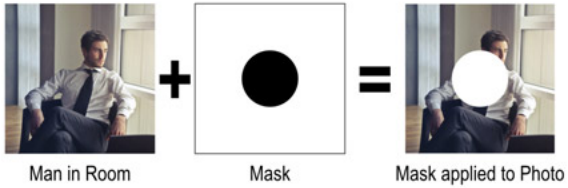

More Complex Masks

Since masks are a special kind of Alpha channel, they can represent values other than 100% active and 0% active (white and black, respectively). Masks can contain values that are intermediate in activity; a pixel in a mask can have a value greater than 0, but also less than 255, and that pixel will be represented as a shade of grey when we visualize the mask.

Masks that contain areas of grey cause the layers to which they are attached to display pixels which are partially transparent. For instance, a mask that has a region of 60% white/40% black will cause the attached photo layer to display pixels in that same region as having an opacity of 60%.

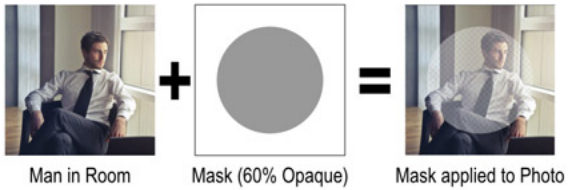

Since the pixels in a mask can contain values from 0 through 255, there is no reason why a mask cannot contain multiple such values. In other words, different pixels in a mask can have different values. If this is the case, then attaching such a mask to a photo would result in different amount of opacity/transparency being applied to different areas of the photo.

Imagine, again, our Man in Room photo with a mask containing many different values. In fact, let's construct a mask using a Gradient – a slow transition from Black to White. This mask contains all values from 0 through 255, and therefore causes the photo to have all degrees of opacity. The photo should be completely transparent where the mask is Black, and completely opaque where the mask is White. The opacity of the photo should fade in from 0% to 100% in a way that corresponds to the transition of Black to White in the Mask.

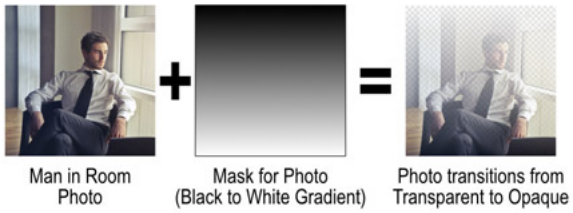

Masks are Not Just for Photos

It turns out that, while masks can change the opacity of images quite nicely, they can also be attached to other parts of your photo edit. Specifically, masks can be attached to Adjustments and Filters – the mathematical constructs that change the color, tone, sharpness, and other qualities of your image.

The idea that an adjustment can have a mask is really very powerful. Because of this, you can

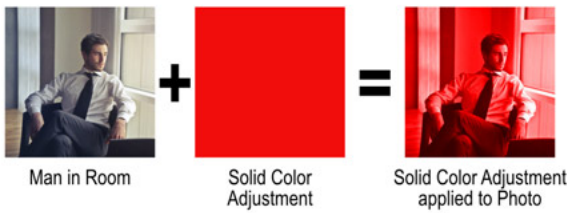

apply an adjustment to your photo, but then limit its effect to only certain areas. When you add a mask to an adjustment, you are telling the adjustment to be opaque (visible) in some areas and transparent in others.

Imagine that we start with our Man in a Room photo, and add an adjustment to change the color to red (called a "Solid Color" adjustment in Photoshop, and a "Recolor" adjustment in Affinity Photo). Without adding any masks, that adjustment process is shown on the previous page.

Attaching a Mask to an Adjustment Layer

However, what if we attach a Mask to the adjustment layer? Remember, we're not adding any alpha channel information to the photo, just to the adjustment. So, we're not introducing any change in opacity to the Man in Room photo; we're changing the opacity of the red Recolor adjustment layer. You can see what this looks like by checking out the graphic in the margin on the right side of this page.

And then take a look at what the layers look like in Affinity Photo. The Recolor adjustment sits on top of the Man in Room layer, and turns that layer red. However, the Circle Mask is a "child" of the Recolor adjustment, adding an alpha channel to the adjustment layer (*not* to the Man in Room photo layer) and alters the opacity of the adjustment.

How do we Create a Mask?

Although there may be a number of very esoteric techniques for creating and editing masks, three basic techniques are used in the vast majority of circumstances.

The first way to create a mask is simply by use of a menu command. You can add a mask to your stack of layers by choosing "Layer > New Mask Layer" in Affinity Photo, or "Layer > Layer mask > Reveal All" in Photoshop. (There are also choices called "New Empty Layer Mask" and "Hide All" in Affinity Photo and Photoshop, respectively, and these will create masks filled with Black instead of White). Once a mask has been created, it can be repositioned in the layer stack, as needed.

A variation on this method is to use the "New Layer Mask" button in either program. This button, found at the bottom of the Layers panel, will create a mask layer within your stack of layers, similar to the way the menu command did.

Once you have created a blank mask, you can paint on it directly (usually with white and black) to create opaque and transparent areas on the layer to which it is attached.

A second method for creating a layer mask, similar to the first but more powerful, is to create a mask based on a selection. If you select an area of your photo first, you can create a mask that incorporates this selection. Seeing this in action is easier than explaining it in words.

Imagine you start with a photo (let's go back to our Man in Room photo) and you select an area within the photo. Once you create your new mask (with the selection active) you will create a mask

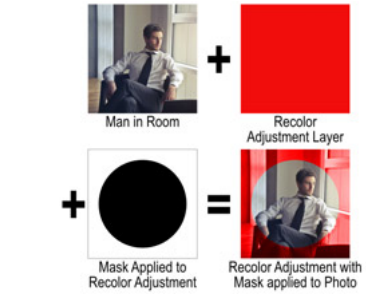

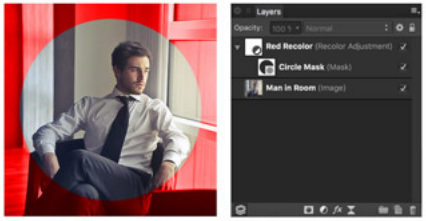

Man in Room with Masked Recolor Adjustment

The Layers as they Appear in Affinity Photo

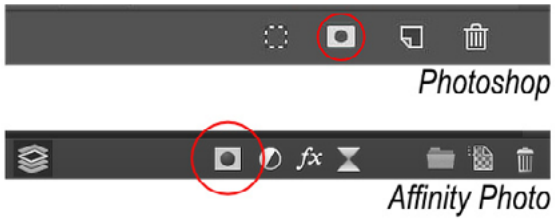

The "Add Mask" Buttons in Photoshop and Affinity Photo

that is already set up to keep the selected area opaque and make the unselected areas transparent. In other words, adding a mask with an active selection creates a mask with the selected area white and the unselected area black. Here's how this looks in editing software (next page).

The third method for creating a layer mask is to transform a pixel layer into a mask. In other words, you can use an existing image to create a mask. Select a pixel layer (it can contain a photo, a hand-drawn graphic, a gradient - anything, really) and use the menus to change the pixel layer into a mask. You can use this mask by attaching it to another layer.

So What's Next?

Now we've introduced color channels and alpha channels, and we've specifically delved into masks – that special case for alpha channels. So far, however, Chapters 1 through 3 have presented facts without much context. The logical next step is to present some real-world situations in which knowing all this stuff might actually make a difference!

In Chapter 4, we'll talk about some of the ways that the color channels can be used when editing photos. Working with only one of the color channels (but not the others) can enhance selections and provide interesting tweaks to color and tonality.

In the final chapter, we'll talk in some detail about how to use masks in a variety of situations. Stay tuned, because the best is yet to come.

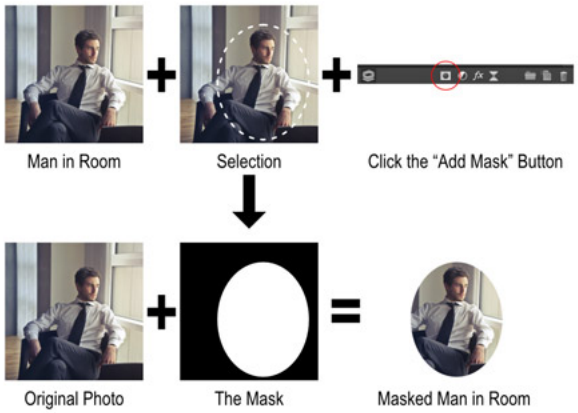

Channel Surfing – Chapter Four

Practical editing using the Red, Green, and Blue color channels

We began this investigation into image channels with an explanation of color channels. Our photos are composed of differing amount of three primary colors – red, green, and blue. Color channels allow us to split a photo into a monochromatic version of each of its primary colors. We can work with an individual color channel in ways that are similar to working with an entire photo. This gives us power to shift, grade, and transpose colors that might not be possible without using individual channels.

Parts 2 and 3 of our journey dealt with alpha channels, and we began to look at selections, opacity, and masks as inter-related forms of these strange creatures.

In Part 4 of this series, we're going to look at some of the ways that color channels can be used in realistic editing scenarios. We will talk about operations on individual color channels, color grading using channels, channel swapping, and (as an introduction to Part 5) using a color channel to aid in creating a selection.

Editing Individual Color Channels

Different software applications will have different methods for isolating a color channel for editing. In Affinity Photo the Channels panel has an eyeball and a pencil icon next to each of a photo's composite channels. These can be used to toggle visibility of the channel, and to choose which channel(s) are being edited. Other software applications access individual color channels in different ways.

As an example, examining the color channels of a sky might reveal it to be quite noisy but, since it is likely to be the red and green channels that are of lower intensity (i.e., “darker” in a monotone representation), the noise might be more prominent there while the blue channel remains relatively free of noise. Apply a noise reduction filter to the red and/or green channels can get rid of noise without affecting any of the clarity of the blue sky.

Some editing adjustments also have the ability to target specific color channels rather than affecting them all equally. A photo of a red flower might benefit from intensifying the reds while diminishing the prominence of the greens and blues.

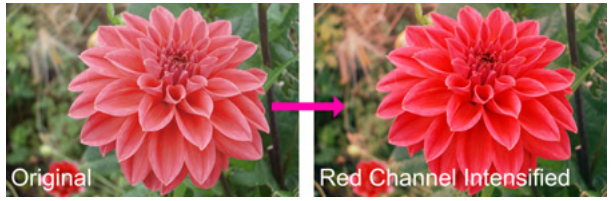

The saturation of the Red Channel has been increased

Color Grading and Balancing

Some editing adjustments can be applied not only to all of the color channels, but can also be applied to individual color channels. In Affinity Photo, the Curves and Levels adjustments are typically used to affect all pixels evenly. However, it is common for these adjustment to offer the

possibility of changing the tonality of the Red, Green, and Blue channels independently of each other. Doing so can affect the color balance of a photo, which can be an effective way to manage color grading and color matching. You might want to check out this YouTube video from PixImperfect that outlines a way to grade skin tones using the Curves adjustment layer.

The HSL (Hue, Saturation, Luminance) adjustment is another adjustment that can not only be applied to all channels, but can target specific color channels one at a time. In our flower example above, the HSL adjustment can be targeted to just the Red channel, and the saturation increased. This leaves the green and blue channels alone, and simply effects changes in the red channel.

The Selective Color adjustment can also change the color balance in specific color channels. Although a bit of an esoteric adjustment, one can target specific color outputs (for instance, the "magenta channel" in the example shown) and alter the input.

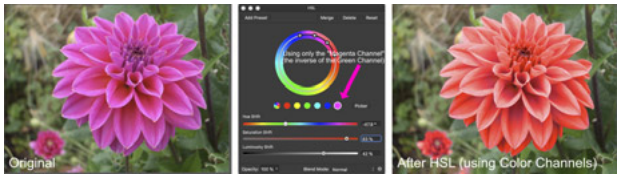

Using HSL to change an individual channel's values

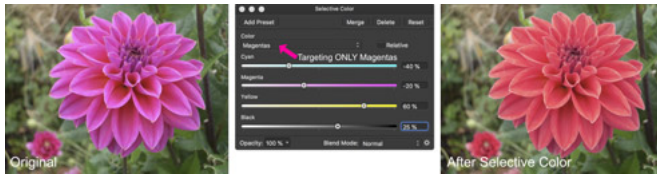

Using Selective Color to alter a color channel's output

Using Curves to adjust tonality of individual channels

Channel Swapping

Channel swapping is a technique used in processing infrared photographs, and in certain other "looks," such as emulating Technicolor processing. When using the channel swapping technique, one channel is mapped to the other and vice versa. For instance, in the Infrared emulation, the Red and Blue channels are swapped. In other words, the values for Red and Blue are switched; all the reds become blue, and all the blues become red.

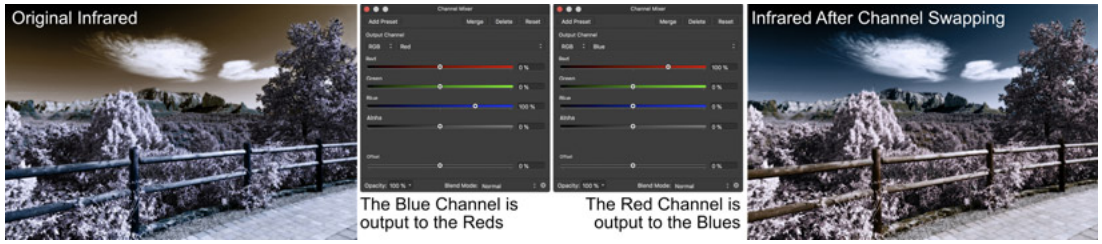

Channel Swapping in a Technicolor emulation

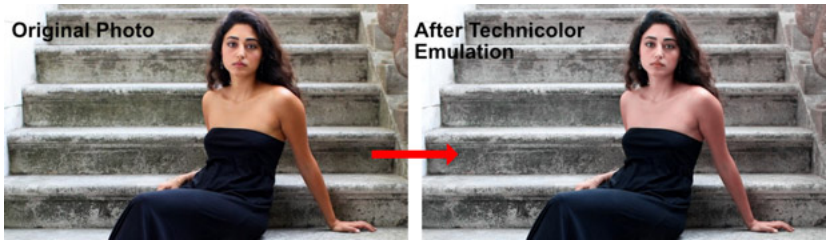

Channel Swapping in an Infrared emulation

Using Channels for Selections

Sometimes a single color channel can form the basis for a selection. This is especially helpful when it comes to making masks, and this starts to cross the boundary between color channels and alpha channels. Particularly when trying to substitute a sky, making an excellent selection can make or break the realism of a mask.

In the example shown, the Blue Channel was isolated, since it forms the best contrast between the sky and the foreground. A greyscale layer was created out of the blue channel, and very light editing was done to make the foreground completely black and the sky completely white. (For those who are interested, a Levels adjustment was made to make the highlights whiter, the shadows darker, and the midtones darker as well. Then, a brush set to Overlay mode painted in black and white to complete the transition to a wholly black and white layer.) Once that was complete, the layer can be converted to a mask. In Affinity Photo, the *"Rasterize to Mask"* command was used; in Photoshop and other applications, the process will be slightly different.

A mask is created starting with the Blue Channel, which makes isolating a selection of the sky much easier

How Do We Finish?

We're almost there. Chapter 5 will explore the use of Alpha channels as a practical part of photo editing. In truth, this is what this series has been leading up to. Alpha channels are where we create masks, and masking is a crucial skill for anyone who wants to do anything beyond the most elementary of photo editing. Bear with me - next month's entry will make it all worthwhile.

Channel Surfing – Chapter Five

Yes, there really are practical uses for Alpha Channels

This is the final chapter of our discussion of image channels. In this article, we will finally look at the practical, everyday use of alpha channels. This final installment has always been our goal; that is, how does the casual photo editor start making use of alpha channels in ways that matter?

We're going to avoid math, we're not going to get too very nerdy, and we're not going to discuss techniques that only three people on the planet would ever use. No, the goal is to make alpha channels accessible to everyone.

And, I'll repeat what I've said before: alpha channels (selections and masks) are the difference between good editing and great editing.

Saving a Selection

We introduced the idea of saving selections in the second part of this series. A selection, as we noted, is really a sort of temporary alpha channel. Some of the pixels in our image are "active" (selected) and some are not. If we were to represent this active/inactive state with a greyscale alpha channel, we would see that the active pixels are white and the inactive (unselected) pixels are black. This is what we would expect, since once part of an image is selected whatever we do next only affects the active pixels.

However, selections are temporary. Once we deselect the selection, that delineation between active and inactive pixels is gone. And, anyone who has labored to make a great selection knows what it's like to accidentally choose "Deselect" and lose all that work.

Photo editing software that lets you manage channels will also allow you to save selections. Doing this creates a new alpha channel. Once you've done this, you can "deselect" and know that your work is safe – the selection is stored as additional information that is more permanent.

Perhaps you've worked hard to select the sky in your photo, in order to help with a sky substitution. But you're not done yet – the selection just isn't perfect. And, it's time for bed, or dinner, or anything else. What do you do? Save the selection as an alpha channel so that you can continue working on it later.

Perhaps you've selected your subject (e.g., the face of the person in your portrait) and you know you're going to want to make some adjustments to it, such as color correction, changing contrast, or adjusting tonality. Save your selection as an alpha channel, and you can use it multiple times for future adjustments.

You might even want to save selections of multiple areas of your photo. The photo on the right has two selections saved – one of just the bird, and one of just a nearby rock. You might want to use

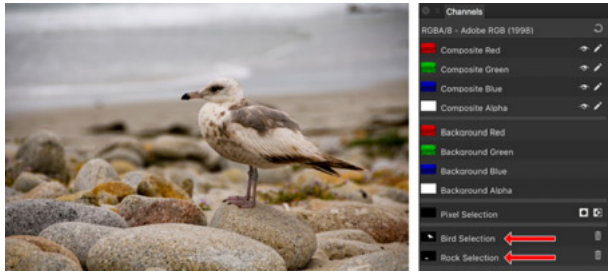

The bird photo has two saved selections – one of the bird and a second of a rock

The rock selectio and the Bird selection (both using Quick Mask for a red overlay)

those selections for masking purposes, or to apply adjustments to those areas individually. You might also want to add those selections together to apply a filter (e.g., sharpening) to just those areas while excluding the rest of the photo.

Compositing Layers

When photographers first start editing photos, they usually learn about masks in the context of merging two pixel layers. In this setting, one can "cut out" part of one photo and "paste" it into another. In general, this is the basis for everything from sky substitutions to the more mundane "put Brad Pitt's face on my boyfriend's body" type of change. Merging two or more pixel layers is known as compositing, since the image which results from it is a composite of the various pictures which went into its creation.

Compositing pixel layers will involve creating a layer stack of 2 or more photos. When you create a mask on the top layer, you can let the bottom layer peek through. So, if I put a landscape with a plain sky on top and a beautiful sky on the bottom I need to create a mask that will hide the sky on the top layer so that my beautiful sky (on the bottom layer) can show through.

The most straightforward way to do this is to work on the top layer, and make a careful selection of the sky. You must be careful to avoid including things like trees and their branches, architectural details on the rooftops, and other things that poke up above the horizon in your photo. Once you have that selection created, first invert it (so that everything BUT the sky is selected) and then click on the "Create a New Mask" button. This creates a "layer mask" attached to the top layer; it is an Alpha channel with the sky showing up as black (since it was not selected and, therefore, not active) and everything else showing up as white.

The result of all this will to hide the area corresponding to the black parts of the mask (the sky) and to reveal the white areas. If you have already placed the more beautiful sky on the bottom of the layer stack, you should be able to see it in the areas that are hidden on the top layer.

Note, also, that you can move the bottom layer around (to position the new sky in the best way) and you can perform adjustments on the foreground and the sky separately, since they are on separate layers.

This type of compositing has the benefit of being non-destructive. Anytime you want to reverse course, you can just turn off the visibility of the mask layer, and prevent the top layer from having any transparent pixels. You can also edit the mask later, in case it was not perfect.

At the right, you'll see the Layers panel for this composited photo (taken from Affinity Photo). The landscape photo is on top, and there is a layer mask attached to it. Below that, there is a second pixel layer called Beautiful Sky. The mask on the top layer allows parts of the lower layer to be seen.

A landscape with a boring sky (left), and a more dramatic sky for substitution (right)

The landscape with its sky masked

The completed sky substitution

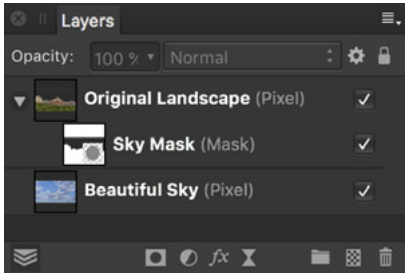

Landscape and Sky layers in Affinity Photo

Applying Adjustments & Filters Locally

Compositing pixel layers is the first step in understanding masks. It is a reasonably straightforward way to visualize just what's going on. It is the digital equivalent of creating a collage.

Once the idea of "revealing and concealing" parts of a pixel layer make sense to you, it's time to expand your horizons. Why not apply the same method to adjustments and filters? After all, if you can recolor, blur, saturate, and change exposure in a photo, why shouldn't you be able to do those things just to a small area of the photo?

Just what does that mean? It means applying a mask to an adjustment layer, or to a live filter (or, in Photoshop talk, a "Smart Filter"). While an adjustment is a simple change to a pixel layer (e.g., make everything brighter; or, make everything green) attaching a mask to an adjustment allows us to limit that change to a selected area.

Maybe seeing this in action will help. At right is a photo of a mom photographing her daughter near the Pacific Ocean. A selection was made of everything BUT the mom and her daughter (i.e., the background only) and an alpha channel was created. The alpha channel was then used to mask (i) an adjustment layer that desaturated and lightened the photo; and (ii) a filter layer that added a Gaussian blur. The end result is that the photo is desaturated, lightened, and blurred, but that these changes were applied only to the background, and not to the subjects.

Sometimes it's easier to use a paint brush to mask an adjustment. In order to do this, we add an adjustment layer and then add a completely black mask to the adjustment. The black mask completely obliterates the visibility of the adjustment (since a black alpha layer renders the layer to which it is attached completely transparent). Then, we paint on the mask with a white paint brush in order to selectively reveal the adjustment only where we want it.

The next set of images shows this process in action. The first image is a photo of a statue. The second image is the same photo, but a blue "recolor" adjustment has been added. In the third image, two things have happened: first, a completely black mask has been applied to the adjustment (which has hidden the blue recoloration completely); and second, a white paint brush has been used to paint on the recolor adjustment's mask. The blue recolor adjustment is revealed in areas where the mask has been rendered white.

The original mom & daughter photo

Alpha channel omits the subjects

Photo with masked adjustments

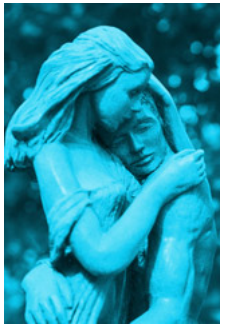

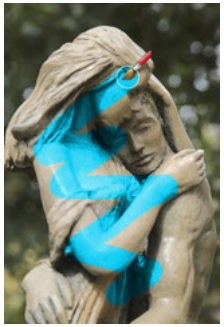

Left: Original Statue Photo
Middle: Recolor adjustment applied to Photo
Right: Adjustment applied locally by using a mask

Masks Based on a Color Selection

Sometimes a mask based on a specific color might be helpful. You might want to pump up the brightness or the saturation of a specific color without affecting other colors. You might want to change one color to another, without affecting the other colors in the photo. You might want to darken, or sharpen, or reduce noise, or any of a dozen other operations, and limit your action to a single color.

You might want to start with the Red, Green, and Blue color channels for your selection. Or, you might want to use whatever variant of "Select Specific Color…" exists in your software. In Affinity Photo, this is called "Select Sampled Color…" and in Photoshop a very powerful panel called, simply, "Color Range" can help select specific colors from your image. A mask based on a color selection helps limit your edits.

Luminosity Masks

Imagine you created a greyscale version of your photograph, and imagine that you turned that greyscale image into a mask. The greyscale version of the photo varies in tonality from black (where the photo is darkest) to grey to white (where the photo is lightest). Once this greyscale image becomes a mask, the light areas are whiter than the dark areas, meaning that the mask allows more visibility of the light areas and less visibility of the dark areas. This is a Luminosity Mask, since it assigns variable opacity of pixels based on their degree of brightness.

If you were to invert that mask, you would have one which favored the darker areas and caused the lighter areas to become increasingly opaque.

Luminosity masks are frequently assigned not to pixel layers but to adjustment layers. In other words, if I want to use a Curves layer to add contrast, but only want to do this to the highlights of my image, I can create a Curves adjustment and use a luminosity mask that favors highlights to limit the adjustment's effects to only the lighter pixels.

Conversely, I might want to increase the exposure of the darker areas of my photo, without affecting the lighter areas. Using an Exposure adjustment and masking it with a luminosity mask that favors the dark areas is one way to accomplish this.

Consider the photo of Bryce Canyon at the right. The highlights are nice, and shouldn't get any brighter. But I'd like to lighten up the darker areas and bring out the details in the shadows. I can apply a Levels adjustment along with an inverted Luminosity mask (that favors the dark areas). This lets my Levels adjustment affect only the dark areas, effectively leaving the light areas alone.

Original image

Levels adjustment applied

Inverted luminosity mask

Levels applied only to Darks

Other Advanced Masks

Any pixels in your photo that can be selected can be turned into an alpha channel. And that alpha channel can be used as a mask.

One of my favorite advanced masking techniques is a Saturation mask. The pixels in a photo are selected based on their saturation; the most saturated colors become the whiter portions of the mask, while the less saturated colors are blacker. The mask that results from a selection based on saturation is a greyscale image, similar in some ways to a luminosity mask. However, a saturation mask is used for different purposes.

Typically, saturation masks are used to control changes in saturation made by adjustment layers (such as a Curves layer or an HSL adjustment). These adjustments can be attached to over- and under-saturated photos.

Applying this mask to an adjustment layer in an over-saturated image lets me lower the saturation more for the already saturated pixels, while protecting the less saturated pixels from going grey. And inverting the mask and using it when adjusting an under-saturated image lets me raise the saturation more for pixels that are less saturated themselves.

Finally, it's Over!

So now our journey is complete. We've explored the Red, Green, and Blue color channels and how they come together to define the color of each pixel. We've introduced the Alpha channel, and seen how it represents opacity and transparency in our photos.

Along the way, we've discussed the concept of opacity in areas of our photos and its editing layers, as well as the nature of selections.

Most importantly, we've tied all this into the concept of masks. Masks are perhaps the most powerful tool in your photo editing bag of tricks. Learn to use them, and reap the rewards.