



# Affinity Photo

from Ten Thousand Feet

*Seth D. Madell*



---

# **Affinity Photo**

## **from Ten Thousand Feet**

**The View from Above**

**by Seth D. Madell**

Copyright 2021 by Seth D. Madell. All Rights Reserved.

Even though this book is provided to you at no charge, the author retains the copyright. You are allowed to distribute the book to others, as long as you do not charge for doing so. Also, you must not make any changes to the book's content, nor represent the content as your own.

Do not create any derivative works that use any part of this book as a portion of your own work.

---

---

# Table of Contents

## First Things First

About the Author.....	2
About this Book .....	3

## Digital Kindergarten

Turning Light into Pixels .....	4
Look Close Enough and We're All Just Pixels .....	5
A Few Words About Resolution .....	5
Raster vs. Vector .....	6

## A Coat of Many Colors

Color Models, Color Spaces, and Color Management .....	8
Channel Surfing.....	14
Alpha Channels – selections, opacity and masks .....	17
Chomping at the Bit.....	20

## The Affinity Photo Paradigm

The Document and the Canvas .....	21
Opening, Saving, Exporting and File Types .....	22
Personas – special case workspaces.....	24
Destructive vs. Non-Destructive.....	25
The Notion of Layers .....	26
Masks .....	29
Layer Opacity, Blend Modes and Blend Ranges .....	33
Adjustments and Filters.....	34
Plug-ins and Macros - extending potential.....	36
The Tools and the Studio .....	36

## What Do I Do Now?

First Steps .....	39
One Final Comment.....	41

## About the Author

I am a retired physician. I tell you this because it also explains what I am not. I am not a journalist or a professional writer. I am not a professional photographer. I am not a software developer. Although I might have a creative streak somewhere inside, I am definitely not an artist. I'm a guy who spent his career trying to explain complicated subjects to people who didn't often understand them. Rather than just tell people "take this pill three times a day," I tried to spend some time explaining how their illnesses and injuries worked, and what we were trying to accomplish. This took a little more time up front, but it paid off as treatment progressed.

So, my goal here is to explain Affinity Photo to the person who doesn't get it just yet. A newbie, so to speak.

I've been taking photos all my life, but only in the past 7-8 years have I approached it with any seriousness. I have learned photo editing on my own, and have certainly had my share of fits and starts along the way. I have always believed that understanding how and why things work will help me use them better; this is the philosophy I bring to this book.

If you are reading this book, you've probably already discovered the Affinity Photo Forum. If you're not already a forum member, I encourage you to join. [You can find the forum at <https://forum.affinity.serif.com>] As of this moment, there are about 140,000 members and a great number of people who peruse the forum as guests. The software forum is a great place to ask questions of other users, and they usually get answered fairly quickly. Forum members may know me by my User Name: "smadell". I have learned a great deal from other members, and have tried to "pay forward" any expertise I have acquired along the way. You are reading one such payment.

## About this Book

I have called this work “Affinity Photo from Ten Thousand Feet” because it tries to give readers an overview of the basics. It deals with concepts rather than details. There are scores of books, tutorials, online videos, and so forth that can explain all the tools, filters, and methods. This is not that book.

In this book, I will try to explain (as simply as I can) some of the stuff you might want to know before you edit your first photo. This is background material.

Some might call this dry and boring. I can live with that. Just remember the old saying: give a man a fish and he’ll eat for a day; teach a man to fish and he’ll eat for a lifetime. This book is about learning to fish.

Here’s what you will not learn. You will not learn what the individual tools in Affinity Photo are for, nor will you learn when to use them. You will not get detailed instructions on any of the adjustments or filters. I will not bore you with the mathematics of each and every blend mode. Most importantly, I am not going to give you step by step instructions on how to achieve a specific look or task within Affinity Photo.

Here’s what you will learn if you continue. You will gain some understanding about digital images in general (pixels, colors, document sizes, and so forth). You will learn what layers are, what kinds of layers exist, and why stacking things on top of one another makes sense. You will learn about transparency (seemingly the buzzword of the age, but for a different reason) and why it means everything in photo editing.

In short, you should start to understand the Affinity Photo paradigm. From there, diving headlong into the various “how to” videos will make so much more sense!

You might also wonder why you were not asked to pay anything for this book. You didn’t have to give me your email address or sign up for a newsletter. You’re not going to receive targeted ads to get you to buy things (at least not from me!) and there is no “Like” button to press. To this, I would say only two things. First, as I said earlier, I am retired. I may not be wealthy, but I don’t need your money. And, I certainly don’t want to have to declare a new business on my tax forms. Second, I suppose I’m a little “old school” and I think that giving things away just feels nice.

Some people say you get what you pay for. Other people say that if you think you got a product for free, then you are actually the product. I hope that neither of these things are true in this case. Enjoy the book. Send me a comment through the forum, if you’d like. And remember to pay it forward.

# Section 1 - Digital Kindergarten

## Turning Light into Pixels

Your camera has a sensor that is sensitive to light. Sensors are described as having a certain number of megapixels and, in general, that means that there are a given number of discrete points on the sensor that pick up data from light. Each of those sensor points is called a *photosite*, and the number of megapixels on the sensor describes how many photosites that sensor contains. For instance, a 24 megapixel sensor has (roughly) 24 million photosites.

When photons hit your camera's sensor, the camera stores information about the intensity of light that struck each particular photosite. So, in the simplest sense, your camera stores a long series of numbers after each exposure – a single number for each of the photosites that may have been exposed to light.

It's actually more complicated than that, since the sensor has red, green, and blue filters that only allow one color of light to register in each of the photosites. So, for instance, a single photosite may only measure the intensity of green light. Red light and blue light are ignored at that photosite. They are measured at the photosites nearby.

The end result of each exposure is a matrix (a "mosaic") of numbers that summarize the intensity of red, green, and blue light at individual photosites.

Each photosite on a sensor corresponds to a single primary color. A single photosite may only supply data about green light, but none about red or blue light. So how can a camera with 24 million photosites create an image with 24 million pixels, where each of the pixels has information about red, green, and blue light? It must seem like somebody is just inventing data somewhere along the way. Put in other terms, if a photosite has information only about green, how come my image has information about red and blue at that same location?

When folks try to explain the difference between Raw files and JPG files, they often say that "raw files don't actually contain images" whereas JPG files do. Of course, that's not entirely true. What is true, however, is that there's an extra step needed for raw files, one that translates that long string of data into usable images. This step is called "demosaicing." It is the method by which data from each photosite (which you'll remember has data about only 1 color) gets information about the other 2 colors it's missing. The process involves interpolation and a lot of guessing, and it's generally pretty accurate, but the conversion from raw data into usable image is that extra step.

Raw files contain only the data from the photosites. The conversion into an image (the

“demosaiicing” step) takes place in your editing software. A JPG file contains data that has already gone through the demosaiicing process (as well as other processes, too).

Once your camera data has gone through the raw conversion steps (demosaiicing and so forth) it is ready to be edited. The image from your camera is now made up of pixels. Let’s talk a bit about them.

## Look Close Enough and We’re All Just Pixels

Pixels are the individual building blocks of an image. Taken by itself, that statement might sound good, but it probably doesn’t mean very much, does it? A pixel is a picture element (that’s where it got its name, after all.) If you put lots and lots of these pixels in a really big grid, and then stand back a little, they start to look like an image.

At its simplest, each photosite from your camera sensor becomes a pixel in your image. Because of that, a pixel has a single color. Because the photosites on your camera sensor are laid out in a grid, so too are the pixels in the images it produces. And while each pixel can have only a single color, if there are enough of them and they’re small enough and closely packed enough, all of those pixels can form complicated and very colorful images.

## A Few Words about Resolution

How big is a pixel? It’s kind of a trick question, because pixels don’t actually have a size. The size of a pixel depends on the size of the image, and it all depends on the physical medium onto which they are printed, or displayed, or otherwise formed. The same pixel can be printed as part of an image on an 8x10 inch piece of photo paper, or it can be projected as part of that same image onto an 80x100 foot billboard. Obviously the individual pixels will be quite small in the first case, and much larger in the second case. Pixel size is an oxymoron until the physical world is used as a display medium.

What does that mean? It means that only when you have chosen the physical size of your image can you determine how large each pixel will become. The size of a pixel is simply a fraction of the size of the image; and the size of the image is up to you. You can display your photos on a 12-inch laptop monitor or on a 27-inch desktop monitor. In each case, all of the pixels can be displayed (within reason, of course). Why? Because the pixels are simply larger on the larger monitor.

Imagine an image that is 100 pixels wide and 80 pixels tall. If that image is printed onto a piece of photo paper that is 10 inches wide and 8 inches tall, each pixel is 1/10 inch square. Now, print that 100x80 pixel image into a poster that is 100 inches wide and 80 inches tall; each pixel

is now 1 inch square.

So, when you buy a new camera do megapixels matter? To a certain degree, yes they do. Think about this not from the perspective of the eventual size of each pixel; think about it from the standpoint of the image you're about to photograph. Point your camera at a beautiful landscape, or a beautiful person, and press the shutter button. If your camera has more megapixels on the sensor, then the eventual image will be made up of more pixels. That means that the pixels are packed together closer, and each particular area of your image is made up of more of them. And, since each pixel only has a single color, having more of them closely packed together means that color transitions will be smoother and details will be crisper.

You might also hear a lot of folks (especially on user forums) asking and/or complaining about dpi and ppi. These stand for *dots per inch* and *pixels per inch*, respectively. They are terms largely borne from the world of printing and publishing, etc. and they attempt to bring together the number of pixels in an image and the size of the image in the real world. After all, if your image has 2000 pixels along one edge, and it is printed on a piece of paper with 10 inches on that edge, the image has 200 pixels per inch. The importance of such a measurement has to do with one's ability to perceive the individual pixels (which is not desirable) in a printed image. It's generally held that printing an image at 200 dpi or greater (some would say 300 dpi or greater) is necessary for images in books, photos in frames, and other images viewed at short distances. Any less densely packed than that and the image lacks good resolution; it looks "pixellated." In short, seeing the individual pixels is not usually the end goal.

Looked at from the opposite viewpoint, the image you get from your camera comes with a fixed number of pixels. For a 24 megapixel camera (which takes photos with a 3:2 aspect ratio) those images will be about 6000 pixels on the long side and 4000 pixels on the short side. If you want to print that picture at no less than 200 pixels per inch, then you can print the picture no larger than 30 inches x 20 inches (still a pretty sizable print).

## Raster vs. Vector

Like many graphics programs, Affinity Photo can handle two very different types of elements. Raster elements are those that are stored as a grid of individual pixels. Perhaps the defining characteristic of a raster element is that each individual pixel can be edited, separate and apart from any of the others. Vector elements are simpler, and are largely defined mathematically. They are shapes, lines, and other simple forms that exist as a unit.

When the user opens a photograph taken with a digital camera, for instance, the image that opens is composed of individual pixels. Each pixel in the photograph can be edited. Each pixel



can be changed, selected, erased, moved, or otherwise altered while the rest of the photograph is untouched.

A vector element can only be edited by changing its globally defined parameters. A vector element (for instance, a rectangle shape) is defined by its size (typically, the coordinates of its left, upper corner and its right, lower corner are all that are needed for this), the thickness and color of its outline (called its “stroke”), and the color of its fill. Changes can be made to any of those larger parameters but, for instance, the pixels which comprise the shape’s fill cannot be individually edited. You cannot erase individual pixels from the shape’s stroke, nor can you copy just a small part of a vector; you can only copy the whole thing.

The benefit of a raster element is that it can be manipulated down to the pixel level. The benefits of vector elements are that they can be manipulated as objects rather than just a collection of pixels, that they are infinitely scalable without losing definition, and that they are simpler and therefore more efficient when computing power is considered.

We will address some of these differences when we discuss the different types of layers that are available in Affinity Photo. For now, the important message is that vector elements cannot be edited at the pixel layer unless they are converted to raster elements. This is done with the **Rasterize** command. Please understand, however, that this is a one-way street; a vector that has been rasterized cannot be turned back into a vector.

## Section 2 - A Coat of Many Colors

### Color Models, Color Spaces and Color Management

Take a deep breath, because the air gets a little rarefied way up here. Color is one of those things that people think they have an intrinsic sense about, only to find that there's so much more to it than they ever could have imagined.

What color is the grass? It's green, sure. And what color is the sky? Blue. Of course. For most of us, our understanding of color stops right about there.

I do not have a PhD in Color Science, nor do I pretend to be an expert on such things. Therefore, this discussion will be profoundly simplistic in some people's eyes. However, judging from the questions that come up on the Affinity Photo user forum very, very frequently, there are many among us who will benefit from even a very simple discussion.


All of you PhD's – you can skip to the next chapter.

#### **Color Models**

What if we had to describe color in a way that made it possible to create a specific color using just a few numbers. If I want to tell you how to recreate a specific color, and I have to do that without showing you the color itself, how would I start?

Color models are mathematical constructs that help us to describe color. We will limit ourselves to the big three (RGB, CMYK, and Lab), but understand that all color models break up colors into specific components, and then quantify how much of each component is present.

The most commonly used color model is RGB. This is the model that our cameras use, and it is the model used in monitors (such as the ones connected to our computers, tablets, phones, and televisions). The common feature among all of these is that they create color by emitting light. All of these devices emit Red, Green, and Blue light. When varying amounts of red, green and blue are emitted in a single area, they combine to form other colors. So, if we want to display the color yellow, we end up displaying equal amounts of red and green light.




The RGB color model is additive, since each component adds to the brightness of the result. That is, it assumes a black baseline and then adds various amount of red, green and blue light. These three colors, shining in varying intensities, combine to create every visible color.

There are two important corollaries to this. First, when there is no red, green, or blue light there is only the black starting point. The absence of RGB is black. Second, when there is maximum red, green and blue light all shining at the same time, this is defined as white. There is actually a third corollary that comes from this, and that is that when equal but intermediate amounts of red, green and blue light shine together, these form shades of grey.

The CMYK color model is a subtractive model, and is used in creating inks, paints, and other physical pigments. This color model is based on each color's amount of Cyan, Magenta, Yellow and Black. The model is called subtractive because increasing the amount of each component makes the end result darker. The CMYK model starts with a baseline of white, and adding components tends toward a maximum result of black.

Perhaps the most important thing to remember, since this book is about photo editing on a computer screen, is that it is almost always the case that one should not be working start-to-finish in CMYK. Most of us should stick to working in RGB, since our end result will either be viewed on a computer screen or other monitor (and these are RGB devices) or they will be printed, in which case the printer almost always assumes that it will be making an RGB-to-CMYK conversion at the time of printing. Working primarily in CMYK is best left for printers and publishers, and they are undoubtedly not the target audience for this type of book.



The third color model that is commonly used is called Lab. Lab color is entirely mathematical, and was created to describe the manner in which colors are perceived by humans. For most of us, Lab color is quite esoteric. The Lab color model uses three components to describe each color: Luminance, the A component (a scale that extends, more or less, from red to green) and the B component (a scale that extends, again more or less, from blue to yellow).

I have found that using Lab color really shines in one respect. Since each color model describes colors based on the amount of the model's components, the Lab model stands out because one of the three components of Lab is "luminance." Lab color separates the lightness of a color from the hue and saturation of a color. Therefore, when a color is expressed in terms of its Lab

color model values, it becomes easier to adjust luminance separately from hue and saturation.

Another thing to understand about color descriptions, irrespective of the color model being used, is that each color is described by a series of numbers. In the RGB model each color is described by a triplet of numbers, each indicating the relative amount of Red, Green, and Blue in that color. When we are working in 8 bit color (more on that in a later section), there are 256 possible amounts of each component. [256 is  $2^8$ , meaning that 8 bits of either 0 or 1 will indicate a possible shade.] Pure red is noted to be: 255,0,0 – 255 is the maximum amount of red, and 0 is the minimum amount of green and blue. Pure white is 255,255,255 and pure black is 0,0,0.

When working in the CMYK color model, each color has four numeric components, since there are 4 different components for each color. The Lab model once again relies on triplets.

### **Color Spaces**

While color models are simply a mathematical way to describe colors, the term “color space” is used to denote a specific set of colors. We should be aware, for instance, that it is impossible to know what the reddest red that is theoretically possible might look like. However, we can define a group of colors in which the “most red” color is well defined.

Sometimes, a group of colors such as this is specific to a device. In other words, a color space might be defined by that group of colors which can be displayed (or printed) by a specific device. At other times, we might create a group of colors that encompass a specific group of hues, levels of saturation, and degrees of lightness. When groups of colors such as this are standardized, it makes for easy comparison. For instance, we might say that your computer monitor can display 100% of the colors in Color Space A, but can only display 90% of the colors in Color Space B.

In theory, I suppose, there are an infinite number of possible color spaces. In general practice, however, there are three that you should know about. These are sRGB, Adobe RGB (1998), and ProPhoto RGB. (There are also common color spaces that are useful when working in CMYK and Lab, but most of us will not be working start-to-finish using any color model other than RGB.)

Without delving into advanced math, three-dimensional models, and esoteric color philosophy, here are a few things you should know about the three most common color spaces.

The most commonly used color space is sRGB. It is the smallest of the three color spaces, meaning that the other two contain colors that are brighter and more saturated than those contained within the sRGB space. There are simply a lot of colors that sRGB does not contain,



and which cannot be displayed within that color space. However, it remains the most commonly used space for one simple reason: most decent monitors can display all (or nearly all) of the colors in that color space.

When sharing documents, posting images online, or otherwise providing graphics to someone else who may or may not be able to display all the colors in a larger color space, sRGB is the color space to use.


Also, there are a lot of applications which will assume that your color images utilize sRGB as their color space and, if they are provided in a different space, they will not display properly.

The second largest color space in common use is Adobe RGB (1998). This color space includes colors which can be brighter and more saturated than sRGB. However, it is commonplace for computer monitors to be unable to display the brightest and most saturated colors.

The largest color space in common use is ProPhoto RGB. (Affinity Photo uses a version of ProPhoto called ROMM. In practical use, they are identical.) ProPhoto contains the largest group of colors, and it contains colors that are brighter and more saturated than any of the other two. However, it is a rare (and very expensive) monitor that is able to display all of the colors in the ProPhoto color space.

So why use a large color space like ProPhoto if computer monitors, many computer applications, and most websites will not be able to make use of it fully? The answer is that it can be preferable to work in a large color space (like Adobe RGB or ProPhoto RGB) but, as a last step, convert the image to sRGB before it is provided to others. This can give you, as an editor, the most leverage in creating color in your document. But, once the document is complete, converting it to a more universally used color space like sRGB makes your color work usable by others.

A quick vocabulary entry is offered. The collection of colors encompassed by a color space is called its “gamut.” A color which falls outside of a color space (usually, one which is brighter and more highly saturated) is said to be “out of gamut.”



## ***Changing Color Spaces***

While we're on the subject, there is a difference between "assigning" a color space to a document, and "converting" it to a different color space. Remember that the colors in your image are stored as a series of numbers. When working in RGB, every pixel has a color described by a triplet of numbers: how much red, how much green, and how much blue. When your document contains pure red, for instance, that red pixel is simply denoted as 255,0,0.

Depending on the color space you are working in, the notation 255,0,0 can translate to different colors! Since ProPhoto has more colors than Adobe RGB, which itself has more colors than sRGB, the color described by 255,0,0 is brighter and more saturated in Adobe RGB than in sRGB, and in ProPhoto RGB it is brighter and more saturated still.

When working in one color space and "converting" to another color space, Affinity Photo must choose a potentially different numeric triplet to describe the same color. So, the color that is 255,0,0 in the sRGB color space will become 219,1,1 when converted to Adobe RGB. It will become 179, 70, 26 when converted to ProPhoto RGB. What this implies (and what is true) is that the values are different because there are brighter and redder colors in those color spaces. The sRGB red might be the brightest and reddest red in the sRGB space, but red can get brighter and redder in larger spaces like Adobe RGB and ProPhoto RGB.

Your software converts colors by storing a new triplet for each color. The goal is to get as close as possible to the same color after the conversion.

Assigning a color space is different. When you assign a color space, you are telling Affinity Photo to keep the existing triplets intact. If the colors shift when you assign a different color space (which they probably will) this is the goal you set. The time to assign a different color space is when the software has chosen the wrong one. If you receive a document which uses the sRGB color space, but your software has assumed it to be in the Adobe RGB color space, you will find that the colors are too bright and too saturated. By assigning the proper color space (instead of converting) you can bring the colors back to where they were intended to be.

Most of the time, you will use color space conversion rather than assignment. Converting to a large color space (like ProPhoto RGB) when you first open your document will give you greater leeway while you are editing. Converting to a smaller color space (like sRGB) at the end of your editing is common, so as to maintain compatibility with other users' software.

## ***Color Management***

The term "color management" is usually used to indicate the process of keeping colors

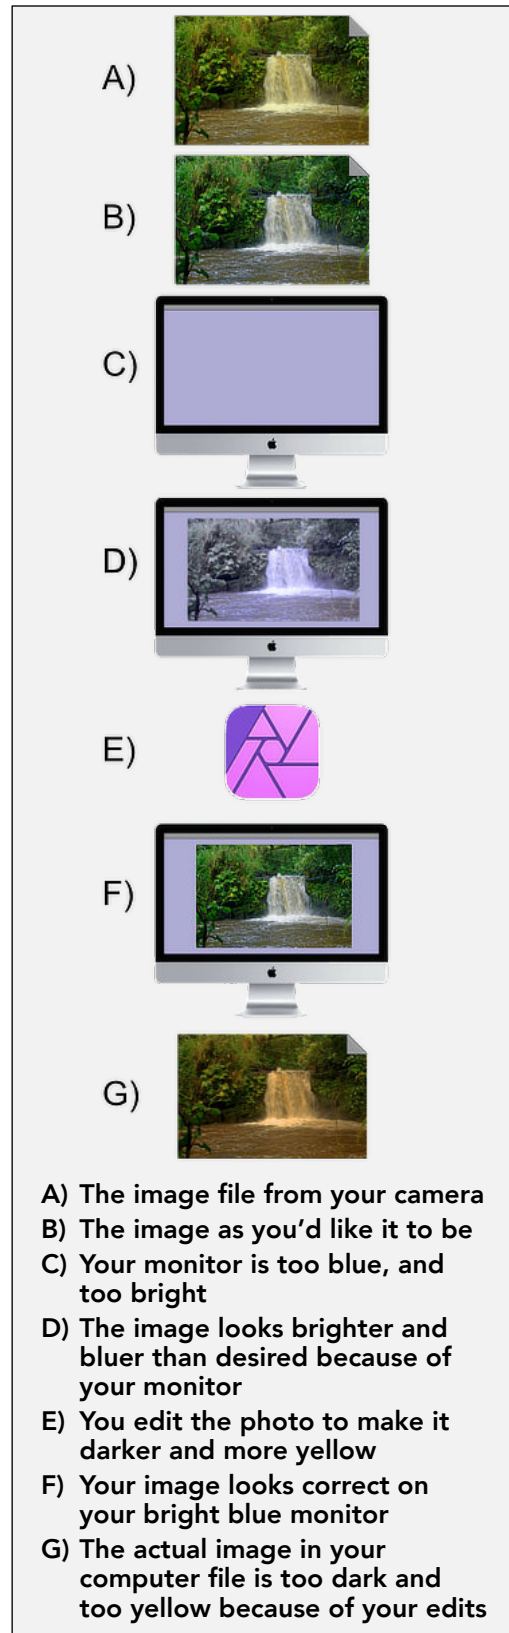
identical across multiple devices. In practical terms, if I create an image with a specific color then I probably want you to see the same color when you open the image file on your own computer.

But, if my monitor is brighter than your monitor, and perhaps has a very blue color cast to the light (this is very common with computer monitors straight from the factory), you might see vastly different colors than I do.

Think of it this way. When I open an image file, the color I see is a combination of the RGB triplets contained in the file plus whatever change(s) my monitor might introduce. So, if my monitor is too bright and too blue, I might edit the photo by adding yellow (to counteract the blue) and decreasing brightness. Unfortunately, if your monitor is less bright and less blue than mine, you might open that photo to see colors that are too yellow and too dark. This is because you are seeing the changes I made, but on a monitor with truer colors. The process is shown in the diagram on the right.

In a similar way, I might get a great result on my monitor only to find that my printer creates photos that are too dark and perhaps the wrong colors. Once again, not only has my monitor introduced changes but my printer can do so, also. (It is also true that my printer prints in a CMYK color space, and that CMYK colors are usually darker and less saturated, anyway.)

The answer to all this is to calibrate the devices in your control. This is done with hardware and software that compares the colors your monitor displays, compared to the colors that the software knows it should expect. If the software sends pure red to the monitor, but it displays a color that is red-with-a blue-tint, then the software knows that the monitor has introduced a specific change. After cycling through a large number of colors, and noting all of the changes made by your monitor, the software creates a *Color Profile*. This is simply a bunch of



mathematical algorithms that counteract your monitor's tendency toward the wrong colors.

Once you create a Color Profile for your monitor, you need to tell your computer system to use that profile for its display. Mac users can do this in the Displays section of the System Preferences app; Windows users will use the Color Management application to set a profile for the Display to use.

The same calibration process can be used when printing at home. Different hardware and software combinations can be used to create a profile for your Printer. Actually, a printer profile should be created for each combination of Printer and Photo Paper you use, since the color of each paper is different and will affect the eventual outcome.

When you purchase photo paper from a manufacturer, it is reasonably common for them to supply printer profiles for each type of paper they sell, and for the more commonly used home photo printers.

When you print your photograph, you will specify the printer profile in the Print Dialog, and this profile will try to counteract any color bias that your printer and paper might be introducing.

The most important thing to remember about display profiles and printer profiles is that they are only to be used by your monitor and your printer, respectively. Do not use a monitor or printer profile as the Color Space for your document. This is an all-too-frequently seen mistake, and it causes a huge amount of confusion when images just don't look right.

## Channel Surfing

We previously talked about Color Models and how each color model described a color by indicating the amount of each of the components it includes. We noted that, for instance, an RGB color is described by a triplet – a series of three numbers, each one varying between 0 and 255 (when dealing with 8 bit color).

Each pixel in our image has a single color, and this color is described using a color model. Usually, as we noted, this will be the RGB model. So, for each pixel there is a corresponding triplet of numbers which describes the 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 border, and y number of pixels from the top border. 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.

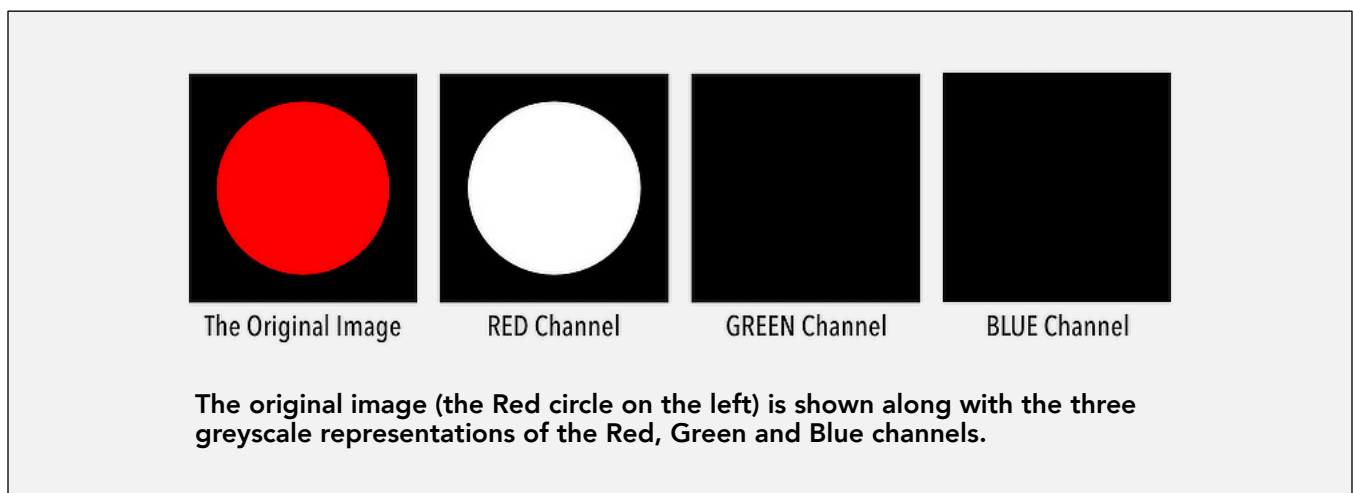


Now imagine, if you will, that our Red pixel spreadsheet (which contains values between 0 and 255 in each of the cells corresponding to the image's pixels) was itself made into an image. Since there's only 1 number in each of the cells, the image could only be a monotone image. In other words, a single "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.

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 below 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 on the following page.



**This new image (on the left) has three circles: one (each) red, green and blue. On the right are the greyscale representations of the Red, Green and Blue Channels. Note how the white areas of the Channel images align with the corresponding colored circle in the original image.**

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).



**In this last example, the image has a circle whose color is made up of varying amounts of red, green and blue. On the right are greyscale representations of the Red, Green and Blue channels, with lighter shades of grey indicating a greater amount of that channel's color in the original image.**

## Alpha Channels – selections, opacity and masks

You might know that an Affinity Photo 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? In the previous section, we introduced color channels and stated that each pixel is composed of red, green and blue components. Different colors simply have higher or lower intensities of these component primary colors. What we didn’t mention is that each pixel has a fourth component, called its Alpha.


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%. In the language of channels, a pixel that is “on” or “active” is represented by white in its Channel representation, indicating that 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.

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 completely inactive.

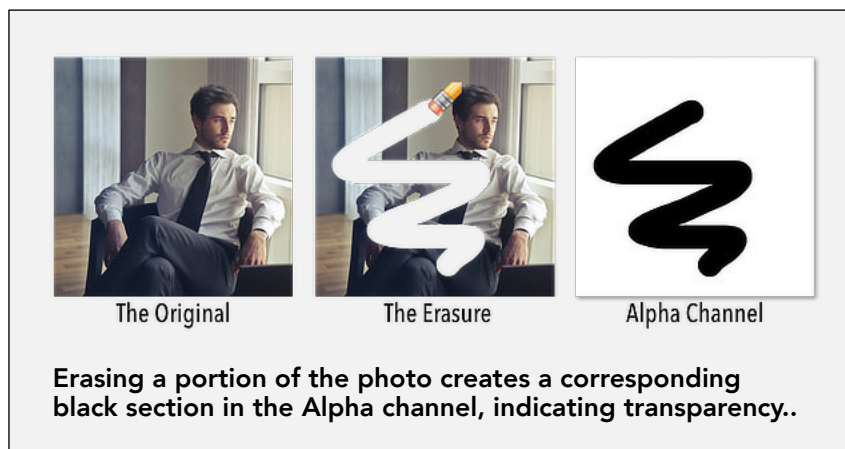
When a layer has an opacity of 60%, its Alpha Channel is 60% white and 40% black (or, just lighter than middle grey) because the opacity setting of the pixels is only 60% active. Here is the Alpha Channel for our Man and Purple Layer stack, where the man has a 60% opacity setting.



### ***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 in the middle of the photo.

If we inspect the Alpha Channel for this image, we'll see two areas distinct from each other. The majority of the photo 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).





### ***Selections are like Temporary Alpha Channels***

Before we go any further, let's take a moment to talk about Selections. You probably know that you can use lots of different tools to make selections. 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.”)

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.

### ***Introducing Masks***

Masks are a special case, and bear some explaining. There will be more in-depth information presented in the following chapter, in its section about Masks. For now, it is enough to understand that masks are the way that transparency can be applied locally. That is, limited areas of an image can be made transparent while other areas are left visible.

Masks involve the creation of Alpha channels. However, instead of an Alpha channel that is all white, all black, or all grey, masks typically have areas that are black as well as areas that are white (and sometimes areas of grey, as well).

Because the Alpha channel has varying areas of black, white and grey, the overall image has varying degrees of opacity. Remember the saying, “Black conceals and White reveals.” That phrase will take you a long way, and we'll come back to it.

## Chomping at the Bit

You might have heard people talking about 8 bit images and 16 bit images (or even 32 bit images). What's going on there?

Once again, assuming that we're discussing an image which uses an RGB color space, 8 bit images allow us to specify 256 different levels (each) of red, green, and blue. In computer-speak, a bit can be a zero or a one – 2 choices. If I were to use 8 bits to describe how much red is in any particular color, then there are  $2^8$  (or 256) different levels of red that are possible. And, if there are 256 possible levels of red, green, and blue (each), then my photo has a palette which contains  $256^3$  possible colors – a total of 16,777,216 different color variations.

While that might sound like way more than you'd ever need, it's not. You will run headlong into the limitations of 8 bit color when you're looking at large areas of similar color. For instance, imagine a photo that has a large area of blue sky. Typically, the sky will be a bit lighter close to the horizon and darker as we look upward. In short, the sky in our photos approximates a gradient that goes from a medium blue to a slightly lighter blue. When using 8 bit color, there are only so many different shades between the blue of the horizon and the blue at higher elevation. When you try to represent this in 8 bits, you will frequently see "banding" – the perceptible stripes of color that occur when that gradient is drawn with fewer shades. Editing photos, especially with adjustments that increase contrast, will tend to exacerbate the problem.

The answer to this is to do your editing in 16 bit color. In a 16 bit image, each color is given 16 bits to specify how much red, how much green, and how much blue is in each particular color. So there are  $2^{16}$  (65,536) variations of red, green, and blue and a resultant palette of about 281 trillion possible colors.

Working in 32 bit color offers even finer shading and smoother transitions, resulting in the absolute minimum amount of banding. Practically speaking, though, for most straightforward editing, using 16 bit color is sufficient.

Even when you start with an 8 bit image, it is generally advisable to convert to 16 bits immediately after opening the image. This will allow your editing to take advantage of the greater number of available color shades while you're working. Even if you save your result as a JPG (which is inherently an 8 bit format) you will see far less banding and generally better results if you do your editing in a 16 bit environment.

## Section 3 - The Affinity Photo Paradigm

This section of the book will deal with concepts that are more specific to Affinity Photo. I will try not to dive too deeply into topics such as (i) how to use each of the tools; (ii) how and when to use any particular adjustment or filter; (iii) how to achieve a specific look by following a pre-defined number of steps. Rather than concentrate on the trees, we will be looking at the forest.

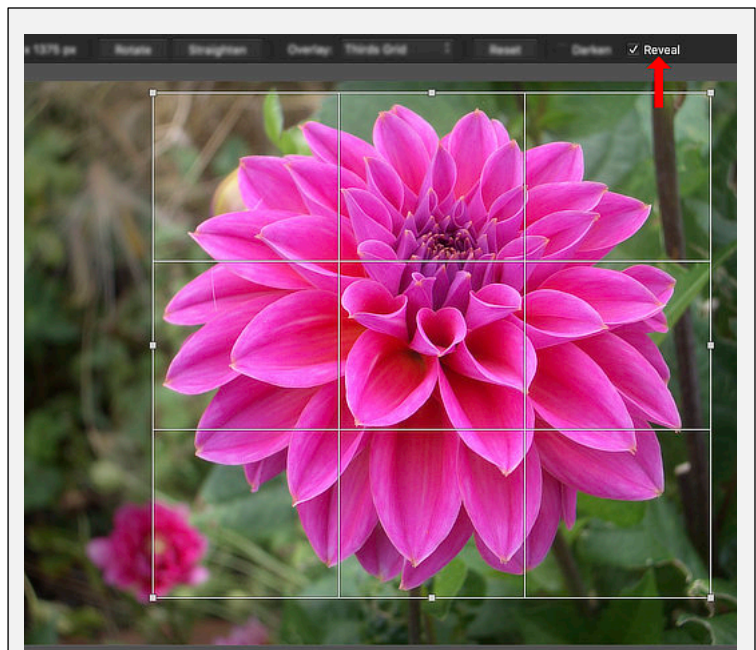
### The Document and the Canvas

Let's first talk briefly about the main window, and the area in which you will be working. You should understand the difference between the Document and the Canvas. Affinity Photo considers the Document to be that part of your work that is visible within the main work area. It is the portion of your work that will be printed (if you issue a **Print** command) and it is the portion of your work that will be exported if you save to a non-native file type (such as creating a JPG file from your image).

The Canvas, however, can contain additional data that lies outside the borders of your Document. Normally, this additional image data remains unseen and is not included in printing, exporting, or other operations. However, it is important to remember that it might be there.

The easiest way to think about the canvas, as opposed to the document, is to (i) open an image (for instance, open a JPG file from your hard drive); and (ii) use the Crop tool to show only a portion of the image. The Document includes only the portion of your image inside the crop. The Canvas, on the other hand, includes the entire image. The portion outside your crop is simply hidden by the borders of your document.

You can visualize the entire canvas most easily by (i) choosing the Crop tool; and (ii) checking the Reveal checkbox. The document should be indicated by the bright white rectangle, while the canvas



The Crop tool displays the contents of the Document (inside the bright white rectangle) as well as the Canvas (which includes image information outside of the bright white rectangle).

includes all areas both inside and outside of the white rectangle.

Sometimes things can stray onto the Canvas, beyond the bounds of the Document, with various processes. These can include copying or placing a second image into your document that is larger than the first document itself; moving one or more parts of your document with the Move tool (the arrow); rotating the image; and various other maneuvers. Always remember that the extra image data is still there unless you explicitly tell Affinity Photo to get rid of it (by using the **Rasterize and Trim...** command).

## Opening, Saving, Exporting and File Types

Computer files exist in a large number of formats. These formats are denoted by their file extensions – the alphanumeric characters at the end of the file name, always following a period. So, for instance, a text file has the extension *.txt* (or sometimes *.text*) while a JPG image file has the extension *.jpg* (or sometimes *.jpeg*).

In general, file formats are of two types. There are application-specific types, usually proprietary and specific to the application that created the file. And there are public, open source file formats which are not specific to any particular application. These more generic file formats are particularly important as they allow various types of data to be shared universally, even by users who do not own the application which created the file.

Affinity Photo uses a proprietary file format denoted *.afphoto*. (At present, Serif – the company which owns and sells Affinity Photo – offers two other applications, called Affinity Designer and Affinity Publisher. These use proprietary formats denoted as *.afdesign* and *.afpub*, respectively.) When you Save a file from Affinity Photo, it is saved into a file format denoted as *.afphoto*. This file format allows Affinity Photo to save all of the data needed to re-open and continue work on your file, without losing any data. (Your *.afphoto* files will often be quite large, in order to maintain this continuing compatibility.)

Affinity Photo can open a variety of file formats. Obviously, it can open its own files and (unless you have changed something in your computer's operating system) double-clicking on an *.afphoto* file will open that file within Affinity Photo. Affinity Photo can also open Affinity Designer and Affinity Publisher files, with complete (or nearly complete) fidelity and no loss of data.

Affinity Photo can open *.psd* and *.psb* files created by Adobe Photoshop, as well as *.ai* files created by Adobe Illustrator and files created by Adobe Freehand. Although these are proprietary file formats, the ability to open and use these files has been around for some time. Affinity Photo can also open many different public file format types. These include DNG, EPS,



GIF, HEIF, JPG, J2K, JPEG-XR, PDF, PNG, RAW, SVG, TGA, TIFF, WEBP, OpenEXR and Radiance HDR.

Note that Affinity Photo can open RAW files, but that most such files (other than *.dng*) are themselves proprietary, and are specific to either camera manufacturers or to specific camera models. Serif maintains an up-to-date list of camera-specific raw files which can be opened by Affinity Photo. Find it at: <https://forum.affinity.serif.com/index.php?/topic/73699-faq-list-of-supported-raw-cameras-19x-for-desktop-lens-correction-list/>

One area of confusion surrounds the file format used within Affinity Photo, while you are working on your image. Once you open a file from your hard drive, Affinity Photo is working on that file in its own format. In essence, a file opened by Affinity Photo (no matter which file type it was before you opened it) now exists as an *.afphoto* file while you are working inside of Affinity Photo. In other words, you can open a JPG or a TIFF file from your hard drive but, once you do, it is effectively an *.afphoto* file while you are working on it.

Another area of confusion surrounds the method by which Affinity Photo saves files. You must understand the difference between the **Save** and **Save As...** commands, and the **Export...** command. Affinity Photo can take your work and create a file on your computer that encapsulates all of that work. If you create an *.afphoto* file based on your work, that file contains all of the data necessary to come back at a later date and pick up right where you left off. However, Affinity Photo can also create other types of files (more on this in a moment) such as JPG files which can be shared with others. Whenever you use the **Save** or **Save As...** commands, you are telling Affinity Photo to create (or update) an *.afphoto* file. In order to create some other type of file (for instance, that JPG) you must use the **Export...** command.

Yes, this is somewhat confusing. Many programs (most notably because of frequent comparison, Adobe Photoshop) can create generic file formats through its **Save As...** dialog. And many folks wish that Serif would have kept to that procedure. But, they didn't. Using the **Save As...** command will always create an *.afphoto* file; use the **Export...** command to create some other type of file.

While we're on the subject, choosing **Export...** allows you to create files with a variety of formats. You can create *.psd* files which can be opened with Adobe Photoshop. You can also create files using public formats such as PNG, JPEG, GIF, TIFF, PDF, SVG, EPS, EXR, HDR, and TGA. Once created, Affinity Photo can be used to re-open those files. However, only the *.afphoto* file format retains all of the information needed to start where you left off inside Affinity Photo. All of the other file formats lose some degree of information during the conversion from Affinity Photo's internal *.afphoto* working format to the public file format you used when exporting.

## Personas – special case workspaces

You might think that an image processing application that you can purchase for under \$50.00 would almost have to be simple and limited in its scope. How could a company like Serif survive otherwise? I don't know how they do it, either, but Affinity Photo packs a lot of functions into its relatively inexpensive package. It has been said that Affinity Photo can do at least 90% (or more) of what Adobe Photoshop can do, and that latter program will cost you far more in the long run.

The folks at Serif decided on a user workspace design that segregated certain parts of the workflow. In many cases, the user can move back and forth freely between these different areas of the interface; in other cases, an activity might be separated from other processes in a “modal” fashion. These various workflow-specific spaces are called “Personas.” There are five main personas, and these can be accessed by clicking on their icons in the Main Toolbar. Most of your work is done in the Photo persona, and many folks will rarely (if ever) leave that persona.

The Liquify persona is used to distort and deform images; the Develop persona is used to develop raw files; the Tone Mapping persona is used for HDR like effects; and the Export persona is used to export files in a more targeted and specific manner than is possible through the **Export...** command.



However, there are other personas available from the File menu, usually accessed when creating a new image. These include: **New Panorama...**, **New Stack...**, **New HDR Merge...**, **New Focus Merge...**, **New Astrophotography Stack...**, and **New Batch Job...** All of these personas open as modal dialogs through which new image files are created and often specific workspaces are accessed (such as the Panorama workspace which is specific to panoramas created through the corresponding menu option).

In the case of the Develop, Liquify, and Tone Mapping personas, it is important to remember that you can bounce in and out of these personas at will during the processing of an image. However, the work you do inside each persona is destructive in the sense that the changes made cannot be undone after leaving that persona.

## Destructive vs Non-destructive

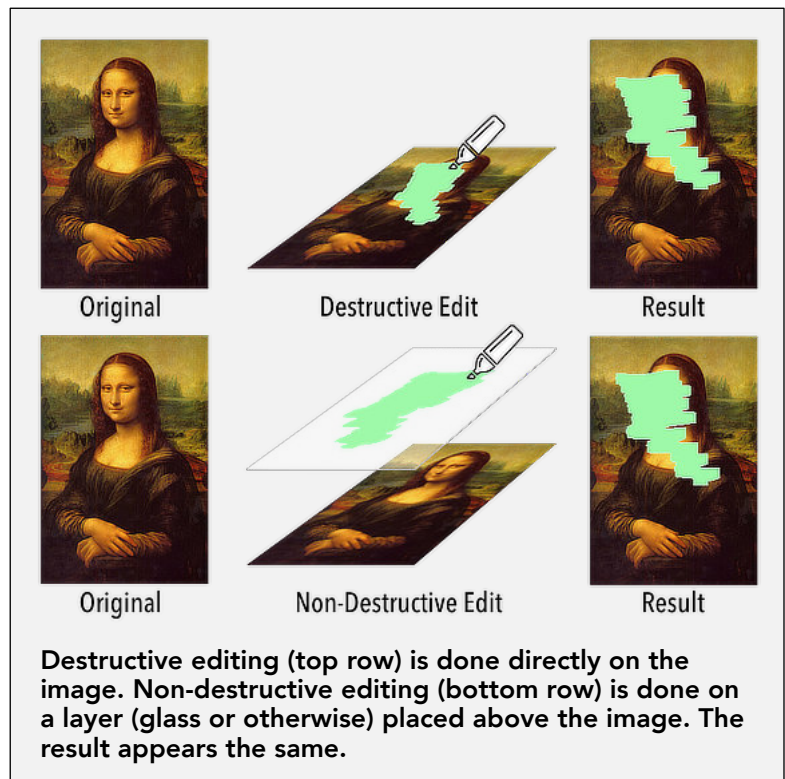
There is a fundamental difference between destructive editing and non-destructive editing. The ability to perform most of your editing in a non-destructive fashion gives Affinity Photo a great deal of its power and flexibility. It also leads directly into the use of Layers, which we will deal with in the next section.

Destructive editing is exactly what it sounds like. The editing that is done to an image forever alters the image. Imagine that you went to your local art museum, found your favorite painting, and drew a line directly on the painting with a wide, green magic marker. Not only would you be in a heap of trouble, but you'd have changed that painting forever. Whatever was once in that area you vandalized, it is now simply a big green blob – forever destroyed by your magic marker transgression.

Now, imagine that instead of drawing your magic marker travesty right onto the painting, you first place a layer of glass directly over the painting and you draw your magic marker line on the glass. To the observer standing in front of that painting, it looks as if you've vandalized the image. There is a great big, wide green magic marker line over a part of the image. But (and here's the important part) the vandalizing was really only done on a piece of glass that is placed in front of the painting; if you remove that layer of glass, no real change was made to the painting itself. This is the nature of non-destructive editing.

In a nutshell, then, non-destructive editing is any change made to an image that can later be hidden, deleted, or further modified. The original image remains untouched and unaltered. The original image has not been destroyed.

The benefits of non-destructive editing should be obvious. If your edits are done in a series of non-destructive steps, it is as if you had done them on multiple layers of glass placed one after



the other above the original image. Not only can you remove a piece of glass (and in doing so hide the editing associated with it) but you can further change that layer of glass (change the parameters of the edit), rearrange the stack of glass layers by moving some in front of others, or you can get rid of all them to bring you right back to the original image – untouched by any of those editing steps.

## The Notion of Layers

When you first open an image (for instance, you open a JPG file created by your camera) you automatically put the pixels from your image into a layer. This first layer, which is automatically named “Background,” represents the original image as it exists in the file on your hard drive. While you could make changes directly to those pixels (paint on them, blur them, erase them, and so forth) those types of changes will permanently alter the pixels. Maybe that’s what you want, but this type of editing removes the opportunity to go back to change an edit, remove an edit (but leave the others) or retrieve the original image. Just as we saw when we drew a green magic marker line on the painting, this is destructive editing and it is (essentially) irreversible.

Imagine, though, that you could digitally put layers of glass on top of your Background layer, performing each of the edits on its own piece of glass. If you looked down on that stack of glass, you would see all of the edits in the same way you saw the magic marker line when you drew it on a piece of glass in front of the painting. However, when you edit in this manner, the editing is non-destructive and reversible.

In Affinity Photo, and actually in many image editing applications, this type of non-destructive editing utilizes Layers. A layer is like a digital piece of glass, and represents a non-permanent change made to the image.

In Affinity Photo, there is no practical limit to the number of layers that can be stacked one upon another. The final image is always formed by looking downward onto the stack from the top. Seen from a different point of view, the original image starts at the bottom of the stack, and is changed incrementally with each successive layer that is placed on top of it. As each layer is added, the image is changed; layers are applied to the image from the bottom up, with the final result being whatever pops out of the top of the layer stack.

Affinity Photo has multiple layer types, and each one has its strengths. As stated at the outset of this book, my purpose is not to tell you when to use one type of layer or another. Nor is it to provide instructions in all the ways each particular layer type is created. However, a basic understanding of what types of layers exist may be helpful as a reference for understanding.

**Pixel Layer:** This type of layer is composed entirely of pixels, and is editable in its entirety. When you first open an image file (for instance, a JPG image) it will open into a layer named “Background” and it will be a pixel layer. Because each and every one of the pixels in a Pixel layer can be edited independently from all the others, you can use any of the Tools directly on the layer. In other words, you can paint, dodge and burn, erase, clone and inpaint, and so forth. If you do this directly onto your Background layer it is a destructive edit and, in general, this should be avoided. However, many of these operations can be performed on pixel layers that are placed above the Background layer, and these edits are by nature non-destructive.

**Image Layer:** If you use the **Place...** command to put a second image into a layer on top of the Background layer, it is put into the Layers stack as an Image layer. An image layer is composed of individual pixels, but they are not individually editable. The image layer exists as a single unit – an object, in computer-speak – and can only be edited by making changes to the entire layer; the layer can be resized, copied, deleted, and so forth, but you cannot paint directly onto the layer, you cannot copy individual pixels from it, and you cannot erase small parts of it. In some ways, it acts like a pixel layer but in other ways it acts like a vector layer (see below). If you want to work on it with the individual tools, or if you want to erase or copy parts of it, you can turn an Image layer into a Pixel layer by using the **Rasterize** command. Understand, though, that Rasterize is a one-way street; there’s no going back.

**Pattern Layer:** Starting with version 1.9, Affinity Photo can create Pattern Layers. These are layers formed first by choosing a size (up to 1024 x 1024 pixels) and then painting directly onto that small area. Whatever you paint into the small area forms the basis of a repeating pattern which fills the entire layer, irrespective of its size. Pattern layers are a great way to create seamless textures and repeating patterns. Pattern layers can be edited just like pixel layers, except that the editing takes place within the original smaller area, and is written into the repeating pattern throughout the layer. Pattern layers can be resized or rotated by transforming the original pixel pattern area. Again, a Pattern layer can be turned into a Pixel layer by using the **Rasterize** command, if needed.

**Fill Layer:** A Fill layer is a way to place a solid color into a layer. This layer can be used to add a color cast or otherwise provide color to your image. The Fill layer usually contains a single color throughout the layer. However, the Fill layer can also contain a Gradient or a Bitmap (in which case the Bitmap can be used similarly to a Pattern layer, in that it repeats through the layer and can be rotated or resized as needed). Fill layers are effectively Vector objects, and cannot be edited on a pixel level.

**Adjustment Layer:** Adjustment layers are those that perform mathematical operations on individual pixels. An adjustment might raise or lower exposure, increase brightness or contrast, change the white balance, and so forth. They are mathematical in nature, and work



on individual pixels (although they generally perform the same math on each and every pixel in an image).

**Live Filter Layer:** Filters are quite similar to adjustments (and, in fact, there are many who note that the differences between them are fairly arbitrary). Filters can be applied destructively by using the Filters menu, or they can be applied non-destructively by using what are known as Live Filter layers. Filters tend to be mathematically more complex than adjustments and use more computational power; because of this they are more likely to slow down processing if you use more than a trivial number of them. Filters often take surrounding pixels into account when making their changes and this accounts, at least in part, for the increase in processing power needed.

**Vector Layer:** Vector layers are created whenever you put a shape, curve, or text into your layers stack. Vectors exist as units, rather than as a collection of pixels, and anything you do to a vector affects the entire layer.

You can create a shape by choosing one of the shape tools found on the Tools panel. These include rectangles, ellipses, stars, arrows, donuts, cogs, and so forth. Drag from top left to bottom right to create a shape. Once the shape has been placed, it can be modified either by dragging the red handles on that shape (not all shapes have these handles) or by setting parameters found on the Context Toolbar.

Text is placed using the Artistic text tool, or the Frame text tool. Text can have a specified font, size, color, stroke, and so forth. Text is also editable, character by character, by selecting the text and using the keyboard (just as you would when working with a word processor).

Curves are more free-form vectors, usually created with the Pen tool. However, you can convert shapes and text to curves by choosing the appropriate button on the context toolbar. Once converted, curves cannot turn back into shapes or text. But, they have all the benefits of curves, which are primarily that they can be edited and changed as to the placement of nodes, curvature of specific segments, joining and breaking of curves, etc.

Curves can also be turned into selections, again by making the appropriate button choice from the context toolbar.

**Group:** Multiple layers can be selected in the Layers panel and placed into a Group. A Group is a special type of layer, which simply indicates that its contents act as a unit. You might want to use Groups for management of the components of a composited image. You might want to put one or more pixel layers into a group with an adjustment, since that adjustment will limit its action to those items within the group.

**Linked/Embedded Document:** You can incorporate external documents into your Affinity Photo layer stack. You might have a JPG or PNG file on your hard drive that you want to bring in as a separate layer. You might have another Affinity Photo document that you want to bring in as a piece of a composite. You can choose **Place...** from the **File** menu and then select the file to bring in. That file become a new layer in your current Document.

Starting in version 1.9, you can specify how such an external file is incorporated into your current document. You can choose a “Placement Policy” on a file by file basis, choosing from either Embedded documents or Linked documents.

Embedded documents are placed into a layer in your current document. They are copied directly into the stack, and they lose any connection with the external document. If that external document is deleted or edited at some point after placing it into your Affinity Photo document, your current layer stack is not affected.

If you choose a Placement Policy that creates Linked documents, then the current Affinity Photo document is dependent on the external document. Any changes you make to the external document can be updated in your current Affinity Photo document. Furthermore, if you delete the external document (or change its name, or move its location) your current Affinity Photo document will lose its link, and your current document may be affected.

**Mask Layer:** Mask layers deserve more discussion, and will be talked about in the following section.

## Masks

We need to talk about masks. Although it is not the purpose of this book to give instructions regarding specific tools or techniques, masks are so integral to photo editing that any conceptual overview must provide some context for their use.

Masks most commonly 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.

### **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)

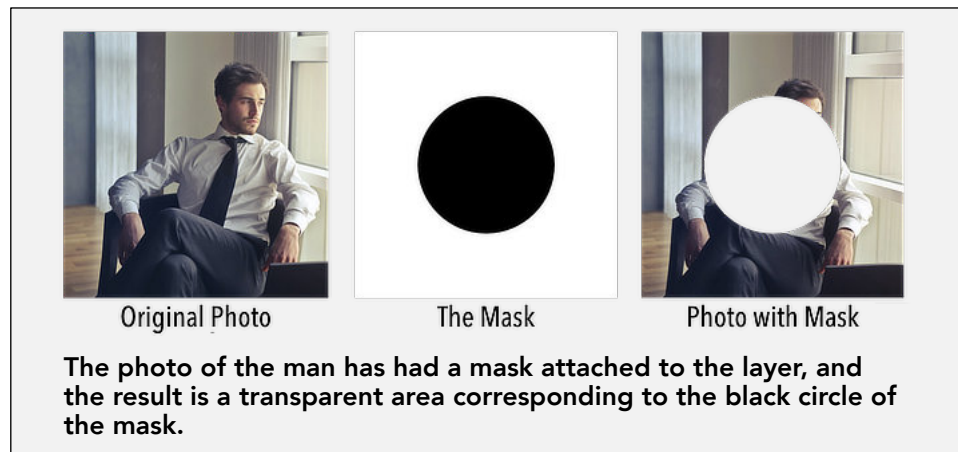
determines the opacity of each of those pixels. Most photos, especially those that have not yet been edited (such as files taken from a digital camera), have alpha channels which are pure white. That is, there are no pixels in the photo which are anything other than 100% opaque.

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.

### ***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.



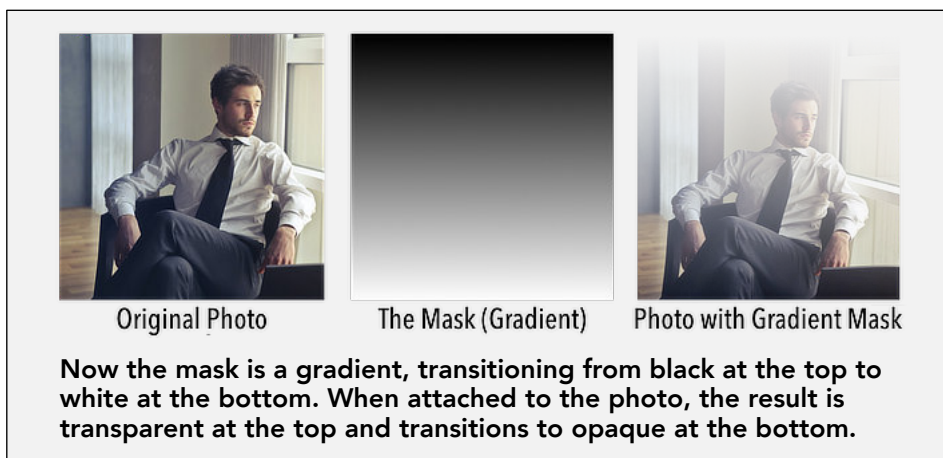
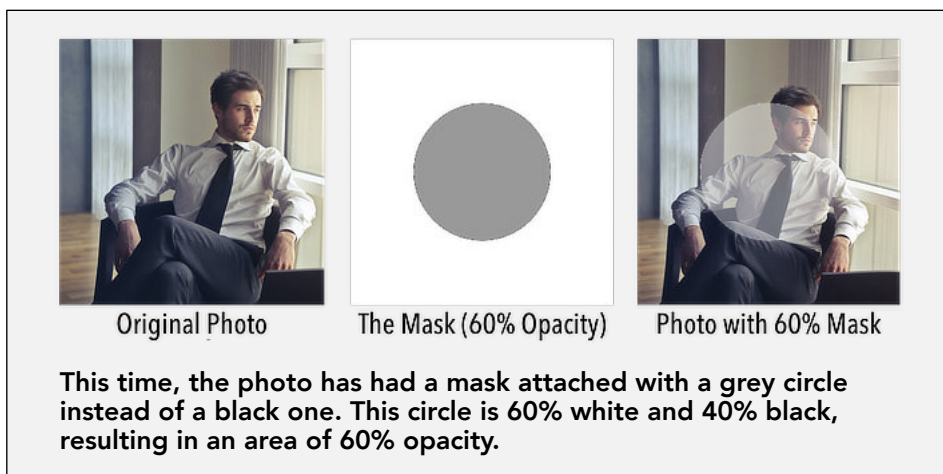
### ***More Complex Masks***

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 used with 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 mask the 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 “Recolor” adjustment in Affinity Photo). Without adding a mask, that adjustment affects the entire image. However, what if we use a mask on the Adjustment layer? If we paint black on the adjustment layer’s mask, then the adjustment is hidden in that area.

In Affinity Photo, it is also helpful to remember that Adjustment layers and Live Filter layers have a mask built in. You don’t have to attach a separate mask layer to an adjustment or a filter; you can simply paint in black, grey, or white directly onto the adjustment/ filter layer and this effectively paints on the layer’s mask.



### **Using Masks Elsewhere**

Masks can be applied to any layer type. So, you can attach a layer mask to an image layer, fill layer, pattern layer, and so forth. Importantly, you can attach a mask to a Group (allowing you, for instance, to group together 2 or more adjustments and then apply those adjustments only to a specified area of the image).

### **Range Masks**

One of the more powerful methods of using masks (and this can apply to masking any of the layer types) is to create a “range mask.” Range masks are a greyscale representation of a layer based on a parameter such as luminosity, color, saturation, and so forth. Perhaps the most common of these is the luminosity mask, which can make a layer more or less transparent based on the luminance of that layer (or of the layer beneath it). Range masks are powerful,



especially when used with adjustments and live filters, since they allow the editor to make an adjustment only to areas of a specific level of luminance (or only to areas of a certain color range, and so forth).

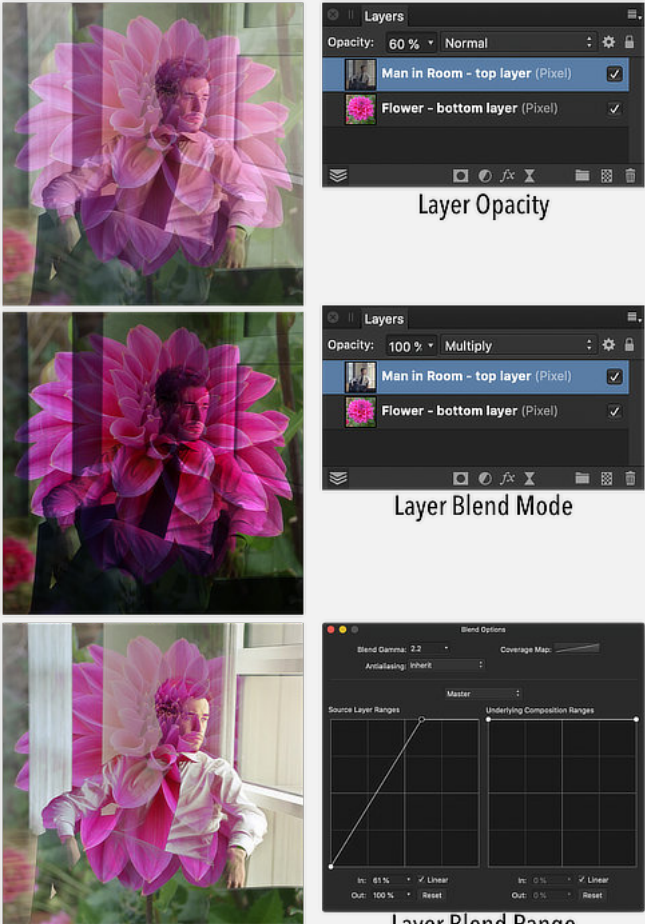
## Layer Opacity, Blend Modes and Blend Ranges

The simplest Affinity Photo document would consist of a single layer. Perhaps you might choose to open a photograph from your camera, make nothing but destructive changes, and then save the result as an edited image. Although this is possible, editing in such a manner means you are missing out on the power of an application like Affinity Photo. If you've read this far into the book, chances are that you're looking for more than that. Assuming that you want to explore Affinity Photo and utilize more of its capabilities, your documents will start to acquire layers – sometimes lots of them.

Once you have at least 2 layers, we need to specify how these layers might interact with each other. For instance, if I open two JPG images and layer one on top of another, the normal state of affairs would dictate that the one on top completely obscures the one on the bottom. Can that be changed, you ask? Of course it can.

### Layer Opacity

Setting the opacity of a layer was something we touched on earlier in the section about Alpha Channels. Each layer in your layers stack has a setting for its opacity. By default, layers start with 100% opacity, meaning that the entire layer is completely visible. Its alpha channel is 100% white. If you change the level of opacity for a level, you make its alpha channel darker until (at 0% opacity) its alpha channel is completely black and the



**Layer Opacity**

**Layer Blend Mode**

**Layer Blend Range**

The Man in Room image is placed on top of the Red Flower image in the Layers stack. Here the layers interact with each other by (top) changing the layer opacity; (middle) changing the top layer's blend mode; and (bottom) setting a blend range for the top layer.

layer itself is completely transparent.

If you were to place one photograph on top of another in the layers stack, lowering the opacity of the top photograph will allow the viewer to see the bottom photo as well.

### ***Blend Mode***

The blend mode for a layer describes a mathematical algorithm for the manner with which it interacts with the layer below it. Without getting too deep in the weeds (I promised I wouldn't do that) blend modes describe the way each pixel in the top layer and the bottom layer are compared and what the result might be.

Each layer added to your layers stack starts with a Blend Mode of "Normal." For pixel, image, and vector based layers, this effectively means that the layer on top obscures the layer below it completely. However, other blend modes change the way in which the top layer interacts with the layer below it. Just as a for instance, the Multiply blend mode darkens the end result based on how dark the top layer is.

### ***Blend Ranges***

Blend ranges are another way to alter the way in which layers interact. Blend ranges alter the opacity of a layer, based on the luminance (the degree of darkness or lightness) either of the layer itself, or of the layer(s) underneath it. As an example, a layer can be told to reveal only the darker pixels, but hide the lighter pixels.

In a similar fashion, Blend Ranges can be applied to an adjustment. The adjustment can be told to apply only to the lightest portions of the layers underneath it, but remain hidden (and therefore not applied) to the darker portions of the underlying layer. This becomes a helpful way of implementing luminosity masking without specifically needing to create a mask layer.

## **Adjustments and Filters**

When editing photos, adjustments and filters will get the lion's share of your efforts. These are the first and most basic steps to take when you want to improve the look of your images. Adjustments affect primarily the color and tonality of pixels, and they tend to apply their changes to each and every pixel of your image. Filters are more computationally intense, and will often distort the pixels of your image; they may induce sharpening, blurring, or other more profound changes.

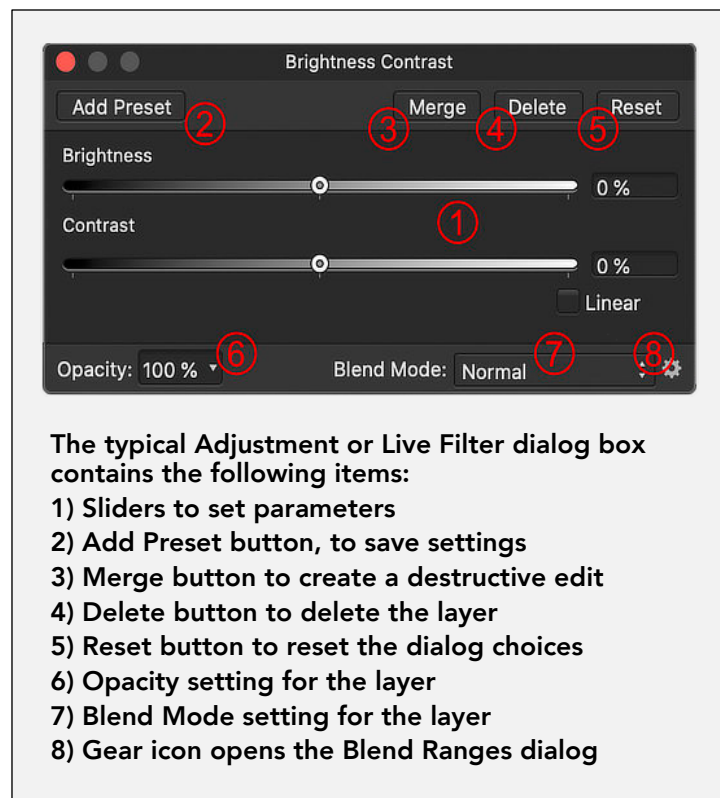
Adjustments can be added by choosing one from the **Layer > New Adjustment Layer >** menu choices. Adjustments include changes to Exposure, Brightness and Contrast, Hue Saturation

and Luminance, Black and White conversion, and many others. Adjustments are always added as a new layer, and are positioned just above the layer currently selected in the Layers panel. (If no layer is selected, then the new adjustment is placed at the top of the layers stack.)

By default, adjustments are non-destructive. Because they are always added as a new layer, they can be hidden, deleted, further edited, or changed in position within the layers stack. You can make an adjustment layer destructive by clicking the Merge button at the top right of the adjustment dialog, and this will write the adjustment's changes into the pixel layer just beneath it as a destructive change.

Filters can be added either as destructive or as non-destructive edits. To edit a layer with a destructive filter, choose the desired filter from the **Filters** menu. Whatever filter you choose will affect only the currently selected layer, and the filter effect will be written destructively into that layer. You can also use a non-destructive form of filter, called a Live Filter layer, by choosing the desired filter from the **Layer > New Live Filter Layer >** menu choices. Not all of the filters are available as live filter layers, but most of them are. Once again, you can add a non-destructive live filter layer, but then Merge the filter into the selected layer by clicking the button in the Filter dialog; this applies the filter destructively, just as if you had chosen it from the **Filters** menu.

No matter which adjustment or filter you may choose, there are a few things which they all have in common. When you first invoke an adjustment or filter (and if you later go back to edit an adjustment or live filter layer) you are presented with a dialog in which you make choices regarding how it edits your image. The example shown on the next page is the Brightness and Contrast adjustment dialog, but its features are typical of all the adjustments and live filters. The dialog box includes sliders (for setting the parameters of the edit), an Add Preset button (which saves the settings for future use), buttons to Merge, Delete, and Reset, as well as choices along the bottom of the dialog which allow the user to choose the Opacity, Blend Mode, and Blend Ranges for the associated layer.



## Plug-ins and Macros - extending potential

The comparison is often made between Affinity Photo and Photoshop. Although the fall-back position has always been to state that “Affinity Photo was not written in order to be a Photoshop clone,” the overlap in function between the two applications is too obvious to avoid making comparisons.

One area where Affinity Photo definitely shows its relative youth is in its ability to extend its architecture beyond that which is already hard-coded into the application. Whereas Photoshop has an extensive third-party collection of plug-ins available, the number of plug-ins that are compatible with Affinity Photo is much smaller (though, to be fair, growing all the time). If you have any familiarity with Photoshop, there is a good chance that your favorite plug-in may not be available for your use within Affinity Photo.

Affinity Photo also lacks an available scripting language. This is something that has been planned for the future but, as of version 1.9, it does not yet exist. Because of this, there is no way to create third-party panels and abilities which appear in the user interface automatically.


Photoshop has a fairly robust ability to create “actions” – multi-step processes that are able to automate specific tasks. These Actions can take advantage of simple programming steps, such as looping, if-then-else steps, user notifications and dialogs, and accepting user input mid-stream. Affinity Photo has the ability to create similar step-by-step processes called “macros.” Macros introduce a great deal of power into the performance of multi-step processes. However, they lack the breadth and depth of Photoshop’s actions; not only is Affinity Photo unable to incorporate a significant number of its operations into Macros, there is no true ability to “program” macros except in a fairly rudimentary way.

Plug-in capabilities and Scripting (and with it, more robust macro creation) are things we can and do hope for in future iterations of Affinity Photo. Right now, however, they are not as comprehensive as a Photoshop user might experience.

## The Tools and the Studio

The Tools are found in a vertical panel along the left side of Affinity Photo’s main workspace. These are individual icons, each of which accesses a means of interacting with your photo. Generally speaking, each tool has its own cursor and performs a unique photo editing function.

The Studio is a term meant to describe the collection of panels included in Affinity Photo that also perform functions in editing your photos. While the Tools are usually constrained to the



The Affinity Photo Studio can include panels in four positions:

- A) The LEFT STUDIO
- B) The RIGHT STUDIO
- C) Docked panels placed off to the side of the workspace
- D) Free floating, single panels situated anywhere within the workspace

left side of the Affinity Photo workspace, Studio panels can be placed in a variety of positions (see the sidebar above). Studio panels can be shown or hidden (by using the **View > Studio >** submenu) and any particular arrangement of studio panels can be saved as a workspace by choosing **View > Studio Presets > Add Preset...** and giving that arrangement of panels a name).

Generally speaking, Tools are used to interact with the currently selected layer. Except in the case of the Text and the Vector tools (shapes and pen-based curves), the tools will not create new layers. (There are exceptions to this rule, however, based on preferences you set up using the Assistant, found on the main Toolbar.)

Often, studio panels can be used to define how each tool is set up. For instance, you might choose the Paint Brush tool to put color onto a pixel layer. However, the color you use is defined in the Color panel – part of the Studio – and the shape and size of the Brush are first specified in the Brushes panel – also a part of the Studio. You might want to add Text to your photo, but the specifics of your font (its size, color, style, and paragraph settings) are all set using various panels from the Studio.



Other panels from the Studio are largely informational, such as the Info and Metadata panels. Others perform functions that tend to be more complex than a single tool could accomplish, such as the Channels panel and the Macro and Library panels.

Other studio panels will create new layers for your image. These include the Layers panel, the Channels panel, the Adjustment panel, the Stock panel and others.

## Section 4 - What Do I Do Now?

### First Steps

#### *Start with some tutorials*

You can't beat video tutorials to teach you the basics. Hearing someone explain a process and watching it happen simultaneously makes the learning process easier. Before you go after video tutorials that teach you to get a specific look or explain a specific effect, find some videos that teach you the very basics.

A good place to start is with the *Affinity Photo Basics*<sup>1</sup> created by Serif – the authors of the application. Find them on YouTube by clicking [here](#)..

You might also want to check out some other free tutorials, including the ones listed below. (These are only videos that I have seen myself, and the list is not meant to be exclusive or otherwise limiting. There are lots of videos out there – paid and free – and this is just a small sampling to help you dip your toes in the water.)

*Olivio Sarikas*<sup>2</sup>: a 17-minute introduction to using Affinity Photo, with good visual examples and well-written, understandable narration. Click [here](#).

*Affinity Revolution*<sup>3</sup>: an old video (from 2016) with basic, good information can be found [here](#).

*InAffinity*<sup>4</sup>: Dave Straker (*dmstraker* on the Affinity Forum) has a YouTube channel with several hundred Affinity Photo videos. They tend to be about 5 minutes long, and vary from the very basic to the most esoteric. Some of them are quite mathematical, and hard for those without computer science background to comprehend. However, Dave also has a series of videos for “absolute beginners” and these are a wonderful place to start a journey. The list of videos is found by clicking [here](#).

#### *Learn to See*

Unless you have a goal in mind, getting a good result is likely to be a matter of luck rather

Copy and paste these URL's if the direct hyperlinks do not work:

<sup>1</sup> Affinity Photo Basics: [https://www.youtube.com/playlist?list=PLjZ7Y0kROWiu\\_QMpLs1twGlFYCdN-1zzu](https://www.youtube.com/playlist?list=PLjZ7Y0kROWiu_QMpLs1twGlFYCdN-1zzu)

<sup>2</sup> Olivio Sarikas: <https://www.youtube.com/watch?v=JaFPoymABvE>

<sup>3</sup> Affinity Revolution: <https://www.youtube.com/watch?v=e8HFGb23wxA>

<sup>4</sup> InAffinity: <https://www.youtube.com/playlist?list=PL0YyTWKOid7E2F02P-sDLvTs3YyE5iAb1>

than purpose. Spend some time staring at a photo when you first open it. I find myself asking, with each new photo, “What is missing?” Is the photo too flat, and lacking contrast? Is it too light or too dark? Would it look better if it were warmer, cooler, or would it be better in black and white? Is there too much there, and should I crop down to include only the most important detail?

Once you have an idea about where you want your image to go, getting there is much easier.

### ***Learn the Adjustment Layers First***

This is where I stray into advice rather than information. With that in mind, the first thing to learn is how to use Adjustment Layers. Adjustments are where the bulk of your editing will take place, especially in the early days. You’ll be able to correct color casts by using the White Balance adjustment. You’ll be able to make images “pop” by using the Brightness and Contrast adjustment, the Levels adjustment, and the Curves adjustment. You can create a wonderful black and white photo using the Black and White adjustment.

Open one of your JPG photos and start adding adjustment layers, one at a time. Play with all the sliders and learn what they do. Try to figure out how one might be different from the other, since this will inform when to use them for a desired effect.

### ***Getting to Carnegie Hall - Practice, Practice, Practice***

There’s an old joke about a guy who’s lost on the streets of New York City. He asks a stranger, “How do I get to Carnegie Hall?” The answer: “Practice, practice, practice.”

Yeah, a time-worn old joke. But with an oversized kernel of truth. The way to learn photo editing is to edit photos! Open some images and try your best. If you like what you see, try to figure out which adjustment or filter or other process got you the result you like. More importantly, perhaps, take stock of the edits that just don’t work.

I believe that we learn more effectively from our failures than from our successes. If you were absolutely sure that a specific edit was going to be the answer to your image’s troubles but it just didn’t turn out that way, spend a few moments trying to figure out why not. It’s one thing to pat yourself on the back for a job well done, and it feels terrific when you can do that; it’s often more instructive, though, to look at your work critically and figure out where a particular set of edits went wrong.

Play. Experiment. Be willing to fail. Most of all, be willing to look at your work with a realistic sense of evaluation. When something just doesn’t feel complete, figure out why. That, in the long run, is the best learning tool available.

### ***Remember to Have Fun***

Unless you're in a very strange profession, no one is going to die if you screw up your photo edit. You can always start over. Photography is supposed to be fun. Even if you approach photography as a career, it's always good advice to "do what you love." If you love photography, it should be enjoyable, right?

When the going gets tough, take a break. Come back to it later. If it takes an extra day to learn a technique, or how to use a specific adjustment, how to write a macro, or anything else, make sure you're enjoying yourself. You can fret over time constraints and accounting once you learn a trade; while you're still in the learning stages, don't forget to have some fun.

### **One Final Comment**

I have been an active member of the Affinity User Forum since 2015, when the application was still in Beta and before it was available for retail purchase. I started simply by lurking and posting an occasional question; at that point, I had little to offer. After a while, I found there were some questions I could answer. Eventually, I posted a few items in the Resources section (mostly macros for other users to download).

Over the years, I have noticed that the same questions tend to come up quite frequently. Many of these deal with background concepts (such as questions about color profiles, bit depth, plug-ins, and so forth) rather than procedural issues (how do I do this...). Writing this book was an attempt to answer some of those conceptual questions.

I do not consider myself an expert in any of this. I have learned enough about these topics to satisfy my own curiosity and to help me get on with using Affinity Photo effectively. I know I'm never going to write my doctoral thesis on any of these subjects. However, at this point creativity is the limiting factor, not my ability to understand the concepts.

I hope that readers have learned something, and that you have found this to be a valuable resource. Please feel free to leave me a comment by way of the User Forum, if you'd like. I'd enjoy hearing from you.

Seth D. Madell  
March, 2021