



Aperture

Digital Photography Fundamentals



 Apple Computer, Inc.
© 2005 Apple Computer, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted for commercial purposes, such as selling copies of this publication or for providing paid for support services. Every effort has been made to ensure that the information in this manual is accurate. Apple is not responsible for printing or clerical errors.

The Apple logo is a trademark of Apple Computer, Inc., registered in the U.S. and other countries. Use of the “keyboard” Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

Apple, the Apple logo, Apple Cinema Display and ColorSync are trademarks of Apple Computer, Inc., registered in the U.S. and other countries.

Aperture is a trademark of Apple Computer, Inc.

Contents

Preface	5	An Introduction to Digital Photography Fundamentals
Chapter 1	7	How Digital Cameras Capture Images
	7	Types of Digital Cameras
	8	Digital Single-Lens Reflex (DSLR)
	9	Digital Rangefinder
	11	Camera Components and Concepts
	11	Lens
	12	Understanding Lens Multiplication with DSLRs
	14	Understanding Digital Zoom
	14	Aperture
	15	Understanding Lens Speed
	16	Shutter
	17	Using Reciprocity to Compose Your Image
	17	Digital Image Sensor
	20	Memory Card
	20	External Flash
	21	Understanding RAW, JPEG, and TIFF
	21	RAW
	21	Why Shoot RAW Files?
	22	JPEG
	22	TIFF
	22	Shooting Tips
	22	Reducing Camera Shake
	23	Minimizing Red-Eye in Your Photos
	25	Reducing Digital Noise
Chapter 2	27	How Digital Images Are Displayed
	27	The Human Eye's Subjective View of Color
	29	Understanding How the Eye Sees Light and Color
	30	Sources of Light
	30	The Color Temperature of Light
	31	How White Balance Establishes Color Temperature

31	Measuring the Intensity of Light
32	Bracketing the Exposure of an Image
33	Understanding How a Digital Image Is Displayed
33	Additive vs. Subtractive Color
34	Understanding Color Gamut
34	Displaying Images Onscreen
35	The Importance of Color Calibrating Your Display
35	Apple Cinema Displays Are Proof Perfect
36	Displaying Images in Print
36	Printer Types

Chapter 3

37	Understanding Resolution
37	Demystifying Resolution
37	Learning About Pixels
38	Learning About Bit Depth
40	How Resolution Measurement Changes from Device to Device
41	Mapping Resolution from Camera to Printer
41	Camera Resolution
42	Display Resolution
42	About the Differences Between CRT and Flat-Panel Display Resolutions
42	Printer Resolution
43	Calculating Color and Understanding Floating Point
43	Learning About Bit Depth and Quantization
44	Learning About the Relationship Between Floating Point and Bit Depth
45	Understanding How Aperture Uses Floating Point

Appendix

47	Credits
----	----------------

An Introduction to Digital Photography Fundamentals

This document explains digital terminology for the professional photographer who is new to computers and digital photography.

Aperture is a powerful digital photography application designed to help you produce the best images possible. However, many factors outside of Aperture can affect the quality of your images. Being mindful of all these factors can help prevent undesirable results.

The following chapters explain how your camera captures a digital image, how images are displayed onscreen and in print, and how cameras, displays, and printers measure image resolution.

How Digital Cameras Capture Images

1

If you've previously shot film and are new to digital media, this chapter is for you. Here you'll find basic information about the types of digital cameras, camera components and concepts, and shooting tips.

People take photographs for many different reasons. Some take pictures for scientific purposes, some shoot to document the world for the media, some make their living shooting products for advertisements, and others shoot for enjoyment or purely artistic purposes. Whatever your reason for picking up a camera and framing an image, an understanding of how cameras work can help you improve the quality of your images.

This chapter covers:

- Types of Digital Cameras (p. 7)
- Camera Components and Concepts (p. 11)
- Understanding RAW, JPEG, and TIFF (p. 21)
- Shooting Tips (p. 22)


Types of Digital Cameras

In its most basic form, a digital camera is a photographic device consisting of a lightproof box with a lens at one end, and a digital image sensor at the other in place of the traditional film plane. Advances in digital photography are fast providing a wide spectrum of features and options that can be challenging for the new digital photographer to master.

There are two basic types of digital cameras: digital single-lens reflex (DSLR) and digital rangefinder.

Digital Single-Lens Reflex (DSLR)

This camera is named for the reflexing mirror that allows you to frame the image through the lens prior to capturing the image. As light passes through the DSLR camera's lens, it falls onto a reflexing mirror and then passes through a prism to the viewfinder. The viewfinder image corresponds to the actual image area. When the picture is taken, the mirror reflexes, or moves up and out of the way, allowing the open shutter to expose the digital image sensor, which captures the image. Most features on a DSLR are adjustable, allowing for greater control over the captured image. Most DSLR cameras also allow the use of interchangeable lenses, meaning you can swap lenses of different focal lengths on the same camera body.




Digital Rangefinder

There are two classes of digital rangefinder cameras: coincident rangefinder and point-and-shoot.

Coincident Rangefinder

Unlike DSLR cameras, the coincident rangefinder does not provide the photographer with the ability to view the subject through the lens. Instead, the coincident rangefinder employs a mirror or prism that uses triangulation to unite the images seen through the viewfinder and a secondary window to bring the subject into focus. The photographer sees two images overlaid on top of one another in the viewfinder, and the image is not in focus until there is a single image. As with DSLRs, most features in a coincident rangefinder are adjustable, allowing for maximum control over the captured image. An advantage to using a coincident rangefinder over a DSLR is that the lack of a reflexing mirror significantly reduces camera shake. Camera shake is due to hand movement or the vibration of the reflexing mirror found in a DSLR, and can cause blurring of the image.




Out of focus
(image overlays not aligned)



In focus
(image overlays aligned)

Digital Point-and-Shoot

This is a lightweight digital camera, aptly named after the two steps required of the photographer to capture an image. Basically, point-and-shoot cameras require pointing the camera and taking the picture without manually adjusting settings such as the aperture, shutter speed, focus, and other settings that professional photographers routinely set on more sophisticated cameras. Of course, some point-and-shoot digital cameras do include adjustable aperture and shutter settings. Point-and-shoot digital cameras are generally light and small, have built-in automatic flash, require no adjusting of focus, and most often include an LCD display that allows you to view the image through the lens in real time via the digital image sensor. Most manufacturers of point-and-shoot cameras separate the viewfinder from the lens assembly to simplify construction and achieve a compact size. The lens, aperture, and shutter are one assembly, irremovable from the camera itself.



Because rangefinder cameras separate the optical path between the viewfinder and the lens assembly, optical compression and frame indicators (guidelines) are used to approximate the image's frame. This approximation often causes subtle differences between what the photographer sees in the viewfinder and what is captured in the image. This is especially noticeable when the subject is close to the camera.


Camera Components and Concepts

The basic components of a DSLR are described below. (Most of the components in a rangefinder are also found in a DSLR.)

- Lens
- Aperture
- Shutter
- Digital image sensor
- Memory card
- External flash

Lens


A lens is a series of sophisticated elements, usually glass, constructed to refract and focus the reflective light from a scene at a specific point—the digital image sensor.



Beyond framing an image, the first interaction you have with the reflective light from your subject is through your camera's lens.

Focal Length

An important attribute of a lens, besides its quality, is its focal length. *Focal length* is technically defined as the distance from the part of the optical path where the light rays converge to the point where the light rays passing through the lens are focused onto the image plane—or the digital image sensor. This distance is usually measured in millimeters. From a practical point of view, focal length can be thought of as the amount of magnification of the lens. The longer the focal length, the more the lens magnifies the scene. In addition to magnification, the focal length determines the perspective and compression of the scene.



Understanding Lens Multiplication with DSLRs

Most interchangeable lenses were originally created and rated for the 35 mm film plane of traditional SLRs. If you compare the area of a 35 mm film plane with the area of most digital image sensors' image planes, you'll see that the area of most digital image sensors is a bit smaller. The focal length of a lens changes when it is put on a DSLR with a digital image sensor smaller than 35 mm. This smaller image plane effectively increases the focal length of the lens because more of the image circle coming out of the lens is cropped. For example, if you put a 100 mm lens on a DSLR that has a 24 mm digital image sensor, the focal length of the lens is multiplied by a factor of approximately 1.3. A 100 mm lens with a 1.3x multiplication factor effectively becomes a 130 mm lens (100 mm multiplied by 1.3).

Another reason to take lens multiplication into account is that shooting wide-angle images becomes increasingly difficult when using cameras with smaller digital image sensors. For example, if your digital image sensor is 24 mm, you require a lens with a focal length less than 24 mm to achieve a wide-angle view. Check your camera specifications for the size of your digital image sensor.

Lens Types

Although there are many varieties of lenses, common lens types include telephoto, wide-angle, zoom, and prime. All of these lenses perform the same basic function: they capture the reflective light from the subject and focus it on the image sensor. However, the way they transmit the light differs.

Note: Although there are several subcategories and hybrids of these lens types, these are the most basic.

Telephoto

A telephoto lens is a lens with a long focal length that magnifies the subject. Telephoto lenses are typically used by sports and nature photographers who shoot their subjects from great distances. Telephoto lenses are also used by photographers who want greater control over limiting the depth of field (the area of an image in focus). The larger aperture settings, combined with the long focal lengths of telephoto lenses, can limit the depth of field to a small area (either the foreground, middle, or background of the image). Small aperture settings, combined with long focal lengths, make objects in the foreground and background seem closer together.

Wide-Angle

A wide-angle lens is a lens with a short focal length that takes in a wide view. Wide-angle lenses are typically used when the subject is in the extreme foreground and the photographer wants the background in focus as well. Traditionally, the focal length of a wide-angle lens is smaller than the image plane. However, in the digital photography age, the sizes of image sensors vary, and the lens multiplication factors of most DSLRs increase the focal length. Check the specifications of your camera to ascertain the size of your digital image sensor. If the size of your digital image sensor is 28 mm, you require a lens with a focal length less than 28 mm to achieve a wide-angle view.

Zoom

A zoom lens, also known as an *optical zoom lens*, has the mechanical capacity to change its focal length. A zoom lens can be extremely convenient, because many zoom lenses can change their focal lengths from wide-angle to standard and from standard to zoom. This eliminates the need to carry and change multiple lenses while shooting a subject or project. However, because of the movement between focal lengths, the f-stops aren't always entirely accurate. To achieve a greater level of accuracy with apertures, many manufacturers have multiple minimum aperture values as the lens moves from a shorter focal length to a longer one. This makes the lens slower at longer focal lengths. (See “[Understanding Lens Speed](#)” on page 15 for an explanation of lens speed.) Plus, a zoom lens requires additional glass elements to correctly focus the light at different focal lengths. It is desirable to have the light pass through the least amount of glass in order to obtain the highest-quality image possible.

Understanding Digital Zoom

The digital zoom feature offered by some camera models does not really zoom in closer to the subject. Digital zoom crops into the center area of the captured frame, effectively enlarging the pixels. This results in a picture with a lower overall image quality. If you don't have a telephoto or optical zoom lens and you want a close-up, physically move closer to the subject, if you can.

Prime


A prime lens, also known as a *fixed lens*, has a fixed focal length that is not modifiable. Prime lenses often have wider maximum apertures, making them faster. For more information about lens speed, see “[Understanding Lens Speed](#)” on page 15. Wider apertures allow for brighter images in low-light situations, as well as greater control over depth of field. Prime lenses are primarily used by portrait photographers. For more information on depth of field, see “[Depth of Field](#)” on page 15.

Aperture

The aperture is the opening in the lens (created by an adjustable iris or diaphragm) that allows light to pass through. The exposure of the image is determined by the combination of shutter speed and the opening of the aperture. The larger the aperture, the more light is allowed to pass through the lens. The aperture is measured in f-stops, and each stop represents a factor of two in the amount of light admitted. The aperture setting (f-stop), combined with the focal length of the lens, determines the depth of field of an image. For more information on depth of field, see “[Depth of Field](#)” on page 15.

f-stop

The photographer adjusts the opening of the aperture by setting the f-stop. An f-stop is a ratio of the focal length of the lens to the diameter of the opening of the aperture. For example, a 50 mm lens with an aperture opened up to a diameter of 12.5 mm results in an f-stop of f4 ($50 \div 12.5 = 4$). Therefore, the larger the numerical value of the f-stop, the smaller the opening of the aperture. The speed of a lens is determined by its largest f-stop value (smallest number). Thus, the larger the aperture, the faster the lens.



Understanding Lens Speed

A lens's speed is determined by the maximum amount of light the lens is capable of transmitting—the largest f-stop value. When a lens is capable of transmitting more light than other lenses of the same focal length, that lens is referred to as *fast*. Fast lenses allow photographers to shoot at higher shutter speeds in low-light conditions. For example, lenses with maximum f-stop values between 1.0 and 2.8 are considered fast.

Depth of Field

Depth of field is the area of the image that appears in focus from foreground to background and is determined by a combination of the opening of the aperture and the focal length of the lens. A small aperture setting results in greater depth of field. Controlling depth of field is one of the easiest ways for a photographer to compose the image. By limiting the depth of field of an image, the photographer can turn the attention of the viewer on the subject in focus. Often, limiting the depth of field of an image helps eliminate clutter in the background. On the other hand, when shooting a landscape, you want the image to have great depth of field. Limiting the depth of field to the foreground would not make sense.

Telephoto lenses (with long focal lengths) tend to have shallow focus when the aperture is opened all the way, limiting the depth of field of an image. Wide-angle lenses (with short focal lengths) tend to create images with great depth of field regardless of the aperture setting.



Shallow depth of field
Only the foreground is in focus.



Great depth of field
The image is in focus from the foreground to the background.

Shutter

The shutter is a complicated mechanism that precisely controls the duration of time that light passing through the lens remains in contact with the digital image sensor. The camera's shutter is activated by the shutter release button.

Prior to the digital age, the shutter remained closed to prevent the film from being exposed. Depending on the type of digital image sensor, a mechanical shutter may not be necessary. Rather than a shutter revealing light to initiate a chemical reaction in the film, the digital image sensor may simply be turned on and off.

Shutter Speed

Shutter speed refers to the amount of time the shutter is open or the digital image sensor is activated. The exposure of the image is determined by the combination of shutter speed and the opening of the aperture. Shutter speeds are displayed as fractions of a second, such as 1/8 or 1/250. Shutter speed increments are similar to aperture settings, as each incremental setting either halves or doubles the time of the previous one. For example, 1/60 of a second is half as much exposure time as 1/30 of a second, but about twice as much as 1/125 of a second.

Photographers often use shutter speeds to convey or freeze motion. A fast-moving object, such as a car, tends to blur when shot with a slow shutter speed like 1/8. On the other hand, a fast shutter speed, such as 1/1000, appears to freeze the blades of a helicopter while it's flying.

Using Reciprocity to Compose Your Image


You can adjust the aperture setting and shutter speed to create several different correctly exposed images. The relationship between the aperture and shutter is known as *reciprocity*. Reciprocity gives the photographer control over the depth of field of the image, which controls the area of the image that remains in focus. This is the easiest way to control what part of the image you want the viewer to pay attention to.

For example, opening the lens aperture by one stop and decreasing the shutter speed by one stop results in the same exposure. Closing the aperture by one stop and increasing the shutter speed by one stop achieves the same exposure as well. Therefore, f4 at 1/90 of a second is equal to f5.6 at 1/45 of a second. The reason is that the camera's aperture setting and shutter speed combine to create the correct exposure of an image.

Digital Image Sensor

When the reflective light from the photographed subject passes through the lens and aperture, the image is captured by the digital image sensor. A digital image sensor is the computer chip inside the camera that consists of millions of individual elements capable of capturing light. The light-sensitive elements transform light energy to voltage values based on the intensity of the light. The voltage values are then converted to digital data by an analog-to-digital converter (ADC) chip. This process is referred to as *analog-to-digital conversion*. The digital numbers corresponding to the voltage values for each element combine to create the tonal and color values of the image.

Each light-sensitive element on a digital image sensor is fitted with either a red, green, or blue filter, corresponding to a color channel in a pixel in the image that is captured. There are roughly twice as many green filters as blue and red to accommodate how the eye perceives color. This color arrangement is also known as the *Bayer pattern color filter array*. (For more information on how the eye perceives color, see “Understanding How the Eye Sees Light and Color” on page 29.) A process known as *color interpolation* is employed to ascertain the additional color values for each element.




Common Types of Digital Image Sensors

There are two types of digital image sensors typically used: a charge-coupled device (CCD) and a complementary metal oxide semiconductor (CMOS).


CCD

CCD sensors were originally developed for video cameras. CCD sensors record the image pixel by pixel and row by row. The voltage information from each element in the row is passed on prior to descending to the next row. Only one row is active at a time. The CCD does not convert the voltage information into digital data itself. Additional circuitry is added to the camera to digitize the voltage information prior to transferring the data to the storage device.



CMOS

CMOS sensors are capable of recording the entire image provided by the light-sensitive elements in parallel (essentially all at once), resulting in a higher rate of data transfer to the storage device. Additional circuitry is added to each individual element to convert the voltage information to digital data. A tiny colored microlens is fitted on each element to increase its ability to interpret the color of light. Advances have been made in recent years in the sensitivity and speed of CMOS sensors, making them the most common type of digital image sensor found in professional DSLRs.



Megapixels

A camera's resolution capability is measured in megapixels. This measurement is based on the number of millions of pixels of image information that can be captured by the light-sensitive elements on the digital image sensor. Thus, a 15 megapixel camera is capable of capturing 15 million pixels of information.

ISO

Traditionally, the International Standards Organization (ISO) has provided a benchmark rating of the relative sensitivity of film. The higher the ISO rating, the more light-sensitive a particular film is. Higher ISO films require less light to record an image. The ISO rating has been redefined for digital cameras, indicating the image sensor's sensitivity to light. Most DSLRs have ISO settings from 100 to 3200 ISO.

Unfortunately, at higher ISO settings (400 ISO and above), some cameras have difficulty maintaining consistent exposure for every single pixel in the image. To increase the sensitivity of the digital image sensor in these situations, the camera amplifies the voltage received from each image sensor element prior to converting the signal to a digital value. As the voltage signals from each element are amplified, so are anomalies within solid dark colors. This results in sporadic pixels with incorrect bright color values, also known as *digital noise*. For more information on digital noise, see "Reducing Digital Noise" on page 25.

Memory Card


After the digital image sensor has captured the image, the camera employs a series of processes to optimize the image. Many of these processes are based on camera settings established by the photographer prior to taking the shot, such as the ISO setting. After image processing, the camera stores the digital information in a file. The type of digital file created varies depending on the camera's manufacturer. However, the camera's RAW file contains the digital image data before it has been converted to a standardized file type, such as JPEG or TIFF. Not all RAW files are alike, but the image data produced by your camera's digital image sensor and processor is retained bit for bit in that file. For more information about these file types, see "Understanding RAW, JPEG, and TIFF" on page 21.

Once the file is ready for storage, the camera transfers the file from its processor to the memory card. There are several types of memory cards, but the process by which they receive the information is the same.

External Flash

Certain photographic situations require the additional light provided by an external flash. Many prosumer DSLR models have built-in or on-camera flashes, but the proximity to the lens and the lack of flash exposure control prevent their use in professional situations.

External flashes provide professional-level control over flash exposure. This allows for fine-tuned fill flash (low-intensity flash that illuminates the subject against a bright background so the subject does not appear in silhouette) and the prevention of overexposed subjects in close-quarter situations.



External or off-camera flashes are synced to the shutter release via the hot-shoe bracket or PC terminal.

Understanding RAW, JPEG, and TIFF

It's important to understand the differences between image file types. RAW, JPEG, and TIFF file types are described below.

RAW

A camera's RAW file is an uninterpreted, bit-for-bit digital image recorded by the camera when the image is captured. Along with the pixels in the image, the RAW file also contains data about how the image was shot, such as the time of day, the exposure settings, and the camera and lens type. This information is also known as *metadata*. RAW refers to the state of the image file before it has been converted to a common format, such as JPEG or TIFF. Because most photography applications previously could not process RAW files, RAW files had to be converted before they could be used in image processing software.

Why Shoot RAW Files?

There are many reasons to capture images as RAW files rather than JPEG files. However, it's important to note that RAW image files require additional work to achieve the color balance you're looking for, whereas JPEG files are color-balanced by the camera for you. JPEG files are also smaller than RAW image files, requiring less storage space.

The advantages to shooting RAW files are:

- Increased bit depth allows for more color-correction "head room." The JPEG format is limited to 8 bits per color channel. RAW images store 16 bits per channel, with 12 to 14 bits per channel of color information. Although it may sound confusing, this means you can do significantly more color correction without degrading the image or introducing color noise. (For more information about bit depth, see "Learning About Bit Depth" on page 38.)
- After the RAW file is decoded, you work with the most accurate and basic data about an image.
- You control the white balance, color interpolation, and gamma correction aspects of the image during post-production rather than when shooting.
- The image file isn't compressed, as JPEG files are, which means that no image data is lost.
- Most cameras are capable of and do shoot color outside the gamut range of JPEG (both Adobe RGB 1998 and sRGB), which means color clipping occurs when you shoot JPEG files. RAW files preserve the camera's original image gamut, allowing Aperture to make image adjustments that take advantage of the full range of captured colors.
- RAW files give you control of noise reduction (luminance and color separation) and sharpening after capture. JPEG noise reduction and sharpening are permanently applied to the image according to the settings on the camera.

JPEG

JPEG (Joint Photographic Experts Group) is a popular image file format that lets you create highly compressed image files. The amount of compression used can be varied. Less compression results in a higher-quality image. When you shoot JPEG images, your camera converts the RAW image file into an 8-bit JPEG file (with 8 bits per color channel) prior to saving it to the memory card. In order to accomplish this, the camera has to compress the image, losing image data in the process. JPEG images are commonly used for online viewing.

TIFF

TIFF (Tag Image File Format) is a widely used bitmapped graphics file format capable of storing 8 or 16 bits per color channel. Like JPEG files, TIFF files are converted from RAW files. If your camera does not have an option to shoot TIFF files, you can shoot RAW files and then convert them to TIFF files using software. TIFF files can have greater bit depths than JPEG files, allowing them to retain more color information. In addition, TIFF files can use lossless compression, meaning that although the file gets a little smaller, no information is lost. The end result is greater image quality. For these reasons, printing is commonly done from TIFF files.

Shooting Tips

Here are some tips for dealing with common photography issues.

Reducing Camera Shake

Camera shake is caused by a combination of the photographer's hand movements or inability to keep the camera still, slow shutter speed, and long focal length. Camera shake results in a blurred image. The focal length of the lens, combined with a slow shutter speed, creates a situation in which the shutter speed is too slow to freeze the image before the camera moves significantly.

You can eliminate camera shake by using a tripod or by increasing the shutter speed to a value higher than the focal length. For example, if you're shooting at a focal length equivalent to 100 mm, you should set your shutter speed to 1/100 of a second or faster. The digital image sensor will capture the image before the movement of the lens has time to register additional light information on the sensor.



Note: Some lenses have image stabilization features that allow the photographer to shoot at a shutter speed whose value is lower than the focal length of the lens.

Minimizing Red-Eye in Your Photos


Red-eye is the phenomenon where people have glowing red eyes in photographs. This is caused by the close proximity of the flash (especially built-in flash) to the camera lens, which causes light from the subject to be reflected directly back at the camera. When the flash fires, the light reflects off the blood in the capillaries in the back of the subject's eyes and back into the camera lens. People with blue eyes are particularly susceptible to the red-eye phenomenon because they have less pigment to absorb the light.



There are a few ways to minimize or eliminate red-eye in your pictures. Some cameras provide a red-eye reduction feature that fires a preflash, forcing the irises in your subject's eyes to close before you take the picture. The main problem with this method is that it often forces subjects to involuntarily close their eyes before the image is taken, and it doesn't always completely eliminate the red-eye effect.


A more effective method is to use an external flash via the camera's hot-shoe mount or, better yet, with an extension bracket. An external flash radically changes the angle of the flash, preventing the lens from capturing the reflection of the blood in the back of your subject's eyes.

While you can also fix the red-eye effect using Aperture, there is no way to accurately reproduce the original color of your subject's eyes. Preventing the problem before it occurs is the preferred solution.



Reducing Digital Noise

Digital noise is the polka-dot effect in images with long exposures or images shot at high ISO settings in low-light situations. The effect is most noticeable in images shot in low-light situations. Many consider digital noise to be a synonym for film grain. Although the causes are the same, the effects are quite different. Some film photographers purposely shoot images with enhanced grain for artistic effect. However, digital noise detracts from the image because of the sporadic bright pixels within solid colors, and lacks the aesthetic qualities of enlarged film grain.



You can reduce digital noise by taking your photographs at ISO settings between 100 and 400. The 400 ISO setting provides more exposure latitude, but even 400 ISO exhibits a little noticeable digital noise. If your subject is not moving and you can't use a flash, using a tripod can allow you to shoot successfully with low ISO settings.

Many DSLR models come with a noise-reduction feature. If you turn on the noise-reduction feature, it is automatically activated when you shoot long exposures. The camera color corrects at the pixel level, processing the image as it's shot. The main negative aspect to digital noise reduction on the camera is the significant lag time required for the image to process between shots. One way to avoid this lag time between shots is to keep the noise-reduction feature on your camera off and use the Aperture Noise Reduction adjustment controls after you've imported your images.

Having a basic understanding of how light is captured, stored, and displayed onscreen and in print can help you achieve the image you intended to create.

It isn't necessary to understand the physics of light and color to appreciate that the colors in an image look realistic. How do you know a sunset is orange, the sky is blue, and the grass green? And exactly how orange is the sunset? What kind of orange is it? It's easy enough to verbally describe your perception of colors, but how do you choose a white balance that conveys the color orange most accurately? This chapter explains how to faithfully reproduce the color you capture with your camera onscreen and in your prints.

This chapter covers:

- The Human Eye's Subjective View of Color (p. 27)
- Understanding How the Eye Sees Light and Color (p. 29)
- Sources of Light (p. 30)
- Understanding How a Digital Image Is Displayed (p. 33)

The Human Eye's Subjective View of Color

Elements of a good photo include composition, color, and brightness. One of your jobs as a photographer is to capture the colors you see as intentionally as possible. Whether you intend to show the color exactly as you see it or you want to enhance the color by adjusting the color temperature, it is your job to understand your choices and intentionally compose your picture.

Unfortunately, human eyes and brains can't be trusted to see colors objectively. Unless you can make side-by-side comparisons of your image on the screen, the photographic print, and the actual subject, it may be hard to tell in what ways the color shifts from one medium to another. Even when making side-by-side comparisons, it is nearly impossible to objectively measure what the differences are when using your eyes alone.

The subjective nature of visual perception should not necessarily be viewed as a handicap. If anything, it may be a blessing. Many challenges in photography come from the fact that the technology is so unforgivingly objective. A common example of this is the issue of white balance. Both film stocks and digital image sensors are designed to interpret white under specific conditions. Outdoor light (daylight) contains a lot more blue light than indoor (incandescent) light bulbs and candlelight. White objects in these different lighting conditions objectively look more blue (daylight), more red (incandescent), or more green (fluorescent), but the brain uses a number of psychological clues to infer that white objects are white, even if they are objectively different.

A white car during sunset objectively looks quite orange, but if someone asks you what color the car is, you would reply with certainty that the car is white. That's because you know the car is white even if it doesn't look white at the moment. In the morning, the car has a bluish tint, and yet again, you would simply say it is white. Digital image sensors and film, on the other hand, record only what they objectively receive, and don't interpret it. The auto white balance feature on many digital cameras measures the scene in the viewfinder and tells the camera to interpret the brightest point as white. This is important to know when switching between different lighting scenarios.

Light and color can be objectively measured and characterized. The scientific analysis of light and color is necessary to build reliable, consistent photographic tools such as film, digital image sensors, displays, and printers. The goal is not necessarily to make all these devices capture or display colors the same way (although this would make things a lot easier), but to develop terminology and processes to objectively measure how these devices are different and adjust output accordingly, so that results match visual perception.

Understanding How the Eye Sees Light and Color

Digital image sensors and the human eye perceive color in similar ways. One of the remarkable things about human vision is the incredible range it has. A healthy eye can see in very bright sunlight and in nearly total darkness. If you have spent much time working with a camera, you know how amazing this range is. Film that works well outdoors is nearly useless indoors, and vice versa. The range of human sight comes from three different parts of the eye:

- *Pupil or iris*: The pupil (also known as the *iris*) contracts and expands depending on the amount of light entering the eye.
- *Rod cells in the retina*: One of the two different types of cells that sense light. Rod cells perceive levels of brightness (but not color) and work best in low light.
- *Cone cells in the retina*: One of the two different types of cells that sense light. Cone cells can perceive color in bright light.

Just as digital image sensors have light-sensitive elements that read red, green, and blue light, the eye has three kinds of cone cells, each sensitive to a different part of the visible electromagnetic spectrum:

- *Cone R*: Perceives colors with red hues with wavelengths in the visible spectrum roughly between 600–700 nanometers (nm).
- *Cone G*: Perceives colors with green hues with wavelengths in the visible spectrum roughly between 500–600 nm.
- *Cone B*: Perceives colors with blue hues with wavelengths in the visible spectrum roughly between 400–500 nm.

The human eye has roughly twice as many green cone cells as red and blue cone cells. This color arrangement is similar to the arrangement of color elements on a digital image sensor. (For more information about how digital image sensors capture images, see “Digital Image Sensor” on page 17.)

The color the eye sees in a scene depends on which cells are stimulated. Blue light, for example, stimulates the blue-sensitive cones, which the brain then interprets as blue. The brain interprets combinations of responses from multiple cones at once and secondary colors are seen as a result. For example, red light and blue light stimulate both the red cones and blue cones, respectively, and the brain interprets this combination as magenta (red + blue). If all three types of cone cells are stimulated by an equal amount of light, the eye sees white or some neutral shade of gray.

Cones are more spread out in the eye than rods. Also, they are much less light-sensitive, so they aren’t even active unless the brightness of a scene or object is beyond a certain threshold. The result is that low-light situations tend to look monochromatic (like black and white), whereas brighter scenes are detected by the cones and thus seen in full color.

Sources of Light

Prior to the invention of electric lights, electromagnetic energy originated from only a few sources. Even today, the sun is the primary source of light. Fire and candlelight provided evening light for thousands of years, though considerably weaker than modern electric lights. Newer sources of light include incandescent light bulbs, fluorescent light tubes, cathode-ray tubes (CRTs), liquid crystal displays (LCDs), light-emitting diodes (LEDs), and some phosphorescent materials. These light sources directly influence the images you create as a photographer.

The Color Temperature of Light

Color temperature is a term used to describe the color of light. Every light source has a color temperature. However, color temperature refers to the color value of the light rather than its heat value. Light’s color temperature is measured in units called *kelvin* (K). This temperature scale measures the relative intensity of red to blue light. Warmer light—light that tends to cast an orangish-red tint across the image—has a lower temperature. Neutral or balanced light occupies the midranges, and has no effect on the image’s color values because of its white qualities. Cooler light—light that is blue in appearance—has a higher temperature.

Light source	Approximate color temperature
Candle	1930 K
Sunlight at dawn	2000 K
Tungsten lamp (incandescent indoor lamp)	2400 K
Typical indoor fluorescent bulb	3000 K
Photographic lamp	3200 K
Photoflood lamp	3400 K
Clear flashbulb	3800 K
Sunlight at noon	5400–5500 K
Blue flashbulb	6000 K
Electronic flashbulb	6000 K
Average daylight	6500 K
Blue sky	12000–18000 K

With the invention of color film came a whole new set of considerations. In addition to correctly exposing the image, photographers had to take into account the various color tints different light sources cast across their film emulsion. Film manufacturers improved the situation by developing film emulsions rated for daylight and tungsten lamp color temperature ranges. Camera manufacturers also jumped in and developed color filters, attached to the camera's lens, to help photographers shoot outside the temperature range of the film. However, these solutions didn't completely eliminate the problem because images shot in unforeseen and adverse lighting conditions remained irreparable during the printing stage.

How White Balance Establishes Color Temperature

When you take a photograph with a digital camera, the color temperature of the scene is not taken into account until the image is processed by the camera's processor. The camera refers to its white balance setting when it processes the image. When the camera's white balance is set to auto, the camera assumes the brightest value is white and adjusts all other colors in the image accordingly. If the brightest value *is* white, the colors in the image are rendered correctly. If the brightest color is yellow, the camera still assumes that value is white, and shifts all the colors out of balance.

However, you can adjust the color temperature of a digital image. White balance is a mathematical process that calculates an image's color temperature and applies the effects to the color values in the image after the RAW image is stored. That color temperature data is stored as metadata in the image. The digital data that makes up the original RAW file is unchanged. So, no matter what white balance or color temperature setting was applied at the time the image was shot, *the color temperature of the image can always be corrected after the fact*. Digital cameras' RAW files solved the problem of color temperature flexibility that the chemistry of film never could.

Measuring the Intensity of Light

In order to shoot an image with the correct exposure, you have to know the correct value of the intensity of light. Photographers use light meters to measure the intensity of the reflective light in a scene. Digital cameras have built-in light meters that are very sophisticated and incredibly accurate. However, their accuracy is subjective. The recommended aperture and shutter values are determined by how light falls in the scene and by how the light meter is set. The camera's light meter may recommend an aperture and shutter combination that offers a decent exposure. However, it may not give you the perfect exposure because it doesn't know what you're photographing. Light meters can't evaluate colors or contrast. They only see luminance, which is the brightness of the reflected light in a scene.

Cameras with sophisticated light meters can be set to meter, or test, specific areas of the scene. Most DSLRs allow you to choose the portion of the viewfinder to meter. These meter settings include, but are not limited to:

- *Evaluative*: Evaluative metering operates by dividing the frame into several small segments, taking a reading from each individual segment, and processing the average of the total segments to recommend the best exposure value for the overall image.
- *Spot*: Spot metering operates by metering within a small target area that is usually in the center of the frame. Spot metering is particularly useful when your subject is placed in front of a relatively bright or dark background. Spot metering ensures that you will correctly expose your subject. The drawback is that the background may be incredibly under- or overexposed. This is why you should bracket (shoot multiple exposures of the same image) when shooting in a situation that requires the use of the spot meter. For more information on bracketing, see “Bracketing the Exposure of an Image,” below.
- *Center-weighted*: When the camera’s light meter is set to center-weighted, the camera measures the light in the entire viewfinder but gives extra emphasis to the center of the frame. This setting is typically used by portrait photographers, because the subject is usually centered and the background isn’t ignored. If the subject moves out of the center of the frame, the meter assumes the background is the correct exposure, leaving your subject incorrectly exposed.

It’s important to point out that light meters provide recommendations only. If the details in the highlights of the scene are more valuable to you, you may choose to expose the image shorter than the light meter recommends. Likewise, if the details in the shadows of the scene are of more value, you may choose to expose the image longer than the light meter recommends. It’s your prerogative as a photographer to use the light meter to obtain the best exposure of the scene in your image.

Bracketing the Exposure of an Image

Even careful metering sometimes yields an under- or overexposed image. This is why professional photographers bracket their images, whenever possible, to be absolutely sure they have a correctly exposed image. Bracketing involves taking three shots of the same image based on the aperture and shutter values recommended by the light meter: one shot underexposed one stop, one shot at the recommended exposure, and one shot overexposed one stop. Shooting the image with a range of three exposure stops is the best way to ensure you’ll have a properly exposed image.


Note: Most DSLR models have a built-in, automatic exposure-bracketing feature. Refer to your owner’s manual for directions about how to use it.

Understanding How a Digital Image Is Displayed

Photographers display their digital images in two basic ways: onscreen or in print. The method by which an image is displayed onscreen and the way it is displayed as a print hanging on a wall are completely different. Computers, televisions, and video and digital still cameras create color images by combining red, green, and blue (RGB) primary colors emitted from a light source. This approach is based on the additive color theory. Printed images require an external light source from which to reflect light. Printing technology uses subtractive color theory, typically with four primary colors: cyan, magenta, yellow, and black.

Additive vs. Subtractive Color

Images with color elements derived from the light source itself are considered to have additive color, while images that subtract or absorb certain wavelengths of light, reflecting back specific colors to the viewer, are considered to have subtractive color. Because of these differences, an image displayed with additive color (for example, on an LCD display) will always look different from the same image displayed with subtractive color (such as on a magazine cover). The reason for this is that digital devices like LCD displays combine red, green, and blue light in different combinations to produce the desired color. All colors combined at their maximum intensities create white, and the absence of color creates black. On the other hand, a printed piece like a magazine cover combines cyan (C), magenta (M), and yellow (Y) inks in different combinations to create a color that reflects the proper color of light. Black ink (K) is added to the image last to generate pure black on the page. The addition of ink creates a darker color, and the absence of ink creates a lighter color. This color process is also known as *CMYK*.



Understanding Color Gamut

In 1931, a group of scientists and intellectuals who called themselves the Commission Internationale de l'Eclairage (CIE) had the goal of defining standards for color. Using as much objectivity as is possible with this highly subjective topic, they developed a coordinate system for categorizing the world of colors. According to this system, every hue the eye can see can be described in terms of x and y coordinates. Taking it one step further, every device that reproduces colors can have its RGB color primaries described by the CIE x and y values. This provides the basis for color-management systems such as ColorSync. The total number of colors described by the two-dimensional plot of these x and y coordinates is often referred to as the device's *color gamut*. In other words, a system's color gamut refers to the total set of possible colors that system is capable of displaying. In addition to this two-dimensional color description, color gamut has a third dimension: its brightness. Unfortunately, the color gamut of displays does not correspond exactly to the subtractive color of print. For example, certain colors that appear onscreen cannot be exactly reproduced in print, and vice versa.

Displaying Images Onscreen

As mentioned earlier, when working with images on your computer screen, you are working with additive light. The display converts electricity into light and the pixels on the screen produce an image by using an RGB color space model. (*Color space* refers to the limits, or parameters, of a given visible spectrum. Common color spaces are sRGB and Apple RGB.) This process begins when the image file on the computer's hard disk is processed and then sent to the graphics card for further processing and temporary storage in memory. The graphics card processes the image, preparing to display it in the specific resolution and color profile of the display or displays connected to the computer. (A color profile is a compilation of data on a specific device's color information, including its gamut, color space, and modes of operation.) Processing the image may take some time, depending on the size and bit depth of the image file, the size and number of displays in the system, and the resolution of the displays. Whether an image was scanned or downloaded directly from a camera, the image was recorded digitally in an RGB color space.

The essence of RGB is the combination of red, green, and blue colors emitted from a light source to form a wide variety of additional colors. On color displays, three colored elements (one red, one green, and one blue) combine to form a pixel. When red, green, and blue are combined at their maximum intensities, the color white is created. When there is an absence of light in all three colored elements, the color black is inferred.

The Importance of Color Calibrating Your Display

It's incredibly important to color calibrate your display or displays to ensure that the color on your screen matches the color you intend to output to print or to the web. Your digital workflow depends on successful color calibration, from capturing to displaying to printing. The adjustments you make to your digital image won't reproduce faithfully in print if your display isn't calibrated. They'll also look different when viewed on other displays. Calibrating your display allows ColorSync to adjust your image for consistent viewing results. Calibrating involves attaching an optical device to your screen that evaluates your screen for luminance and color temperature. There are several companies that manufacture color-calibration tools. The tools can be expensive and can vary greatly in quality, so make sure you do an adequate amount of research before you make your purchase. For a list of available color-calibration tools and devices, see the Mac Products Guide at <http://guide.apple.com>.

Apple Cinema Displays Are Proof Perfect

Apple Cinema Displays are so good at displaying color that you can use them in a SWOP-certified soft-proofing workflow. Display-based proofing systems Remote Director 2.0 from Integrated Color Solutions, Inc. and Matchprint Virtual Proofing System-LCD from Kodak Polychrome Graphics both have Specifications for Web Offset Publications (SWOP) certification. The prestigious SWOP certification means you can use Remote Director 2.0 to approve jobs for press production onscreen without the need for paper proofs, providing significant time and cost savings for print professionals.

Certified systems are capable of producing proofs visually identical to the SWOP Certified Press Proof as defined in ANSI CGATS TR 001, Graphic Technology. Integrated Color Solutions, Inc. and Kodak Polychrome Graphics chose Apple flat-panel displays because they are capable of providing the luminance and color gamut necessary to create an onscreen proof that has the same brightness and feel as paper.

Note: Your Apple Cinema Displays must be color-calibrated to achieve accurate results when soft-proofing your images.

Displaying Images in Print

Displaying images in print requires converting the color from the RGB color space to CMYK. The reason for this is that printed images need to reflect light from external light sources to be viewed. Images are usually printed on white paper, so no white ink is necessary. Darker colors are created by adding colors together, whereas lighter colors are produced by reducing the color mix.

For additional information about image quality in print, see Chapter 3, “Understanding Resolution,” on page 37.

Printer Types

The following printer types are divided into two groups: personal printers and professional printers.

Personal Printers

There are two basic types of printers that are affordable for most photographers.

- *Inkjet*: Inkjet printers create images by spraying little ink droplets onto the paper. Inkjet printers are capable of placing the microscopic droplets on the paper with great precision, resulting in high-resolution photographs. There are two methods of applying the ink to the paper. One technique involves heating the ink to a temperature warm enough to allow the ink to drip. The second method involves vibrating a tiny valve filled with ink, forcing it to fling a droplet onto the page.
- *Dye sublimation*: Dye sublimation printers create images by heating colored ribbon to a gaseous state, bonding the ink to the paper. The ribbon is a plastic material that makes the print nearly waterproof and difficult to tear. The incredible durability of dye sublimation prints gives them a longevity that cannot be surpassed by any other medium.

The quality of inkjet printers has improved remarkably in the past few years, making their resolution and color gamut superior to those of dye sublimation printers.

Professional Printers

There are two basic types of printers employed for professional use. Unlike personal printers, these printers are relatively expensive.

- *Offset press*: Offset presses are used for high-volume printing for items such as magazines and brochures. Offset printing presses deposit ink in lines of halftone dots to produce images on the page. The printer uses a fixed drum to roll the image onto the paper.
- *RA-4*: RA-4 printers are capable of printing digital files on traditional photographic paper. They use a series of colored lights to expose the paper, which blends the colors together to produce continuous-tone prints. Due to their expense and size, most photo-direct printers are only available at professional photo labs.

The concept of resolution often confuses people. Cameras, displays, and printers measure resolution in different ways.

Resolution describes how much detail an image can hold. This section explains image resolution and shows how understanding image resolution can help you create better digital images.

This chapter covers:

- Demystifying Resolution (p. 37)
- How Resolution Measurement Changes from Device to Device (p. 40)
- Mapping Resolution from Camera to Printer (p. 41)
- Calculating Color and Understanding Floating Point (p. 43)

Demystifying Resolution

An image's resolution is determined by the image's pixel count and the bit depth of each pixel.

Learning About Pixels

A pixel is the smallest discernible element in an image. Each pixel displays one color. A pixel's color and brightness range is determined by its bit depth. For more information, see "Learning About Bit Depth" on page 38.

Pixels are grouped together to create the illusion of an image. On color displays, three color elements (one red, one green, and one blue) combine to form a pixel. As the number of pixels increases, the image's detail becomes sharper, more clearly representing the original subject. Therefore, the higher the pixel count, the more likely the displayed image will look like the original subject.

Because so many pixels fit in even a small image, pixel count is often expressed in megapixels (millions of pixels). For example, 1,500,000 pixels equals 1.5 megapixels.

Learning About Bit Depth

Bit depth describes the number of tonal values or shades of a color each channel in a pixel is capable of displaying. Increasing the bit depth of color channels in an image's pixels exponentially increases the number of colors each pixel can express.


The initial bit depth of an image is controlled by your camera. Many cameras offer several file settings; for example, DSLR cameras usually have two settings, allowing the photographer to shoot an 8-bit JPEG file (with 8 bits per color channel) or a 16-bit RAW image file (with 12 to 14 bits per color channel).

Image file types use static bit depths. JPEG, RAW, and TIFF all have different bit depths. As you can see in the table below, the file type you shoot your images in dramatically impacts the tones visible in your images.

Bit depth per color channel	Possible tonal values per color channel	Nearest equivalent file type
2	4	
4	16	
8	256	JPEG, some TIFF
12	4096	Most RAW
14	16,384	Some RAW
16	65,536	Some TIFF


Note: The bit depth of an image file is uniform (each pixel in the image has the same number of bits) and is initially determined according to how the image was captured.

Here's a practical example of bit depth. To understand the effect of bit depth on an image, look at the picture of the girl below, which is an 8-bit grayscale image. Her eye is used to illustrate the effects that lower bit depths have on the resolution of the image.




Formats like JPEG use 24 bits per pixel: 8 bits for the red channel, 8 bits for the green channel, and 8 bits for the blue channel. An 8-bit color channel can represent 256 possible values (2^8), while three 8-bit color channels can represent 16,777,216 values (2^{24}). RAW image files also use three color channels. Because most RAW files have the capacity to capture 12 to 14 bits per color channel, their range of colors is exponentially larger.

The following example illustrates how increasing the bit depth of a pixel increases the number of color values it can represent. Increasing the bit depth by 1 bit doubles the number of possible color values.



How Resolution Measurement Changes from Device to Device

As you now understand, resolution in itself isn't complicated; it simply measures how much detail an image can hold. However, as resolution is described for different digital devices—cameras, displays, and printers—the different units of measurement can be confusing. A camera's resolution is calculated by the number of megapixels (millions of pixels) its digital image sensor is capable of capturing. A display's resolution is expressed in pixels per inch (ppi) or as a maximum dimension, such as 1920 x 1280 pixels. A printer's maximum resolution is expressed in dots per inch (dpi)—the number of dots it can place within a square inch of paper. These changing units make it hard to keep track of the resolution of your digital image as it moves from one device to another. Not only do the units of measurement change, but the numerical values change as well.



Mapping Resolution from Camera to Printer

Tracking the changing units of measurement from camera to display to printer is confusing. But without an understanding of how resolution changes between devices, you can inadvertently compromise the quality of your images.

Camera Resolution

A camera’s potential resolution is measured in megapixels (the number of millions of pixels used to record the image). The larger the number of megapixels, the more information is stored in the image. The reason a camera has a potential resolution is that lens quality, the ISO setting, and the compression setting can affect the quality of your image. For more information on how a camera operates, see Chapter 1, “How Digital Cameras Capture Images,” on page 7.

The number of megapixels a camera is capable of capturing can be used to roughly determine the largest high-quality print that the camera is ultimately capable of producing.

Megapixels	Print dimensions at 200 dpi	Approximate uncompressed file size
1	4" x 3"	1 MB
1	4" x 3.5"	2 MB
2	6" x 4"	3 MB
2.5	10" x 6"	7 MB
4	12" x 8"	12 MB
5	14" x 9"	15 MB
7	16" x 11"	21 MB

Display Resolution

The maximum number of pixels that can appear on a display's screen determines its maximum resolution. Most displays have a variety of resolution settings from which to choose. For example, the 23-inch Apple Cinema HD Display has resolution settings from a minimum of 640 x 480 to a maximum of 1920 x 1200 pixels. As a photographer, you will want to operate your display at its maximum resolution setting. This ensures that you see as much of the image as possible on your screen.

About the Differences Between CRT and Flat-Panel Display Resolutions

CRTs and flat-panel displays are not bound by the same resolution characteristics. CRT displays are capable of resolution switching, so that the resolution you select is displayed at the actual resolution, and the pixels are drawn properly and sharply at any supported resolution. Flat-panel displays have only a single native resolution that appears sharp and true, which is the maximum resolution. Choosing any other resolution forces the entire screen image to be interpolated to that size, resulting in a soft, or slightly blurred, image.


Printer Resolution

In the end, it's the quality of the print that counts. The quality of the print is determined by the combination of two factors:

- *Image file resolution:* The resolution of the image file is determined by the number of pixels in the image and the bit depth of the pixels themselves. Obviously, the more pixels the image file has, the more information it's capable of displaying. However, along with the number of pixels, the bit depth plays a large part as well. The greater the bit depth, the more colors a pixel is capable of displaying.

For more information on bit depth, see “[Learning About Bit Depth](#)” on page 38.

- *Printer resolution:* A printer's resolution is determined by how closely together it is capable of placing dots on paper within a square inch, measured in dpi. A printer's maximum dpi value determines the highest-quality image it can print.



Calculating Color and Understanding Floating Point

As you've learned, digital devices translate color into numbers. Aperture calculates color using floating point, a type of calculation that allows calculations to be performed at a very high resolution with a minimum of error.

Learning About Bit Depth and Quantization


When you capture an image using a digital image sensor, the analog voltage values have to be converted to digital values that can be processed and then stored. For more information, see “Digital Image Sensor” on page 17. The process of converting an analog voltage value to a digital value is known as *digitization*. In the process of converting an analog voltage value to a digital representation, *quantization* must be performed, converting the values to discrete numerical values. The accuracy of each pixel's value is determined by the length of the binary word, or bit depth. For example, a 1-bit binary word can represent only two possible states: 0 or 1. A 1-bit system cannot capture any subtlety because no matter what the tonal value is, a 1-bit system can represent it either as 0 or 1 (off or on). A 2-bit binary word can represent four possible states: 00, 01, 10, or 11. And so on. Most digital RAW image files capture a minimum of 12 bits per color channel (4096 possible states), allowing for many subtle degrees of tonal values to be represented. The more bits available for each sample, the more accurately each color channel's tonal value can represent the original analog voltage value.

For example, suppose you use 128 numbers to represent the tonal values of color channels in each pixel in an image within a range of 1 volt. This means your camera's analog-to-digital converter is precise to 1/128 of a volt. Any subtle variations in tonal values that are more detailed than 1/128 of a volt cannot be represented, and are rounded to the nearest 1/128 of a volt. These rounding errors are known as *quantization errors*. The more the signal is rounded, the worse the quality of the image.

Learning About the Relationship Between Floating Point and Bit Depth

When you make multiple adjustments to a digital image, the adjustments are mathematically calculated to create the result. Just as with analog-to-digital conversions, there can be quantization errors when adjustments are calculated. For example, consider the following calculation: $3 \div 2 = 1.5$. Note that for the answer to be accurate, a decimal point had to be added for an extra level of precision. However, if the bit depth of your pixels does not allow this level of precision, the answer would have to be rounded to either 2 or 1. In either direction, this causes a quantization error. This is particularly noticeable when you try to return to the original value. Without the precision of floating point, you're left with $1 \times 2 = 2$ or $2 \times 2 = 4$. Neither calculation is capable of returning the original value of 3. As you can see, this can become problematic when adjustments require a series of calculations and each subsequent value is inaccurate. Since a large number of calculations are required to perform complicated adjustments to an image, it is important that the adjustments are calculated at a significantly higher resolution than the input or output resolution in order to ensure the final rounded numbers are more accurate.

In the example below, a green channel of a 24-bit pixel (with 8 bits per color channel) is capable of displaying 256 shades of green. If an adjustment is made calling for a calculation between the 167th and 168th color values, without floating point the application would have to round to one or the other. The result of the final calculation would be a color that is close but not accurate. Unfortunately, information is lost.



Understanding How Aperture Uses Floating Point

Internally, Aperture uses floating-point calculations to minimize quantization errors when image adjustments are processed. Floating-point calculations can represent an enormous range of values with very high precision, so when adjustments are applied to an image, the resulting pixel values are as accurate as possible. Often, multiple adjustments to an image create colors outside the gamut of the current working color space. In fact, some adjustments are calculated in different color spaces. Floating point permits color calculations that preserve, in an intermediate color space, the colors that would otherwise be clipped.

When it's time to print the image, the output file has to be within the gamut range of the printer. A pixel's tonal values can be processed with incredible accuracy and then rounded to the output bit depth, whether onscreen or print, as necessary. The accuracy is most noticeable when rendering the darker shades and shadows of the image. The bottom line is that image processing using floating-point calculations helps produce extremely high image quality.

For more information about color gamut, see “Understanding Color Gamut” on page 34.

Credits



Photography by Norbert Wu (pages 41 and 43)

Copyright 2005 Norbert Wu

<http://www.norbertwu.com>



Photography by Matthew Birdsell (pages 9 and 16)

Copyright 2005 Matthew Birdsell

<http://www.matthewbirdsell.com>