TECHNICAL GUIDE TO THE SCIENCE AND TECHNOLOGY OF ELECTRONIC IMAGING

BY
PH2 GEORGE SISTING
US NAVY RESERVES

Index

Index

TOPIC **PAGE**

Camera System

AD Converter

In a digital camera's image system, the sensor device (CCD / CMOS) outputs an analog signal (a voltage) which is then amplified and processed by an Analog to Digital Converter (ADC). The ADC value represents (after processing) the amount of light which was captured by a particular photosite (pixel) on the sensor. Most consumer digital cameras use 8-bit ADC's. This means they may have up to 256 distinct values for the brightness of a single pixel, newer prosumer and most Digital SLR's have 10 or 12-bit ADC's which enable them to encode that same brightness level using 1024 or 4096 distinct values.

Often tied closely together. The ADC does not affect the sensors dynamic range, which is defined by the "depth of the well" (number of electrons a single photosite can hold), a 10 or 12-bit ADC simply allows for a finer graduation of the differences between one brightness and the next. Obviously, sensors with considerably higher dynamic range (such as found in Kodak's Pro DCS, Nikon D1, Fujifilm S1 Pro and Canon D30) will gain some advantage from 10 or 12-bits of data otherwise their dynamic range would be "compressed" to fit into the 8-bits of data. This is a reason cameras such as these utilize a RAW format which records 10 or 12-bits of data per pixel.

AF Assist Lamp

Some manufacturers fit their cameras with a lamp (normally located beside the lens barrel) which is used to illuminate the focus subject in low light. This lamp assists the camera's focusing system where other cameras autofocus will likely have failed. These lamps usually only work over a relatively short range, up to about 4 meters.

Some lamps use Infrared light instead of visible light which is better for "candid" shots where you don't want to startle the subject. Notable higher end flash systems (Canon's 550EX and Nikon's SB-28DX) feature their own assist lamps with far greater range.

AF Servo

Autofocus Servo means the camera's ability to continuously focus on a moving subject. This isn't something that most consumer grade digital cameras can do but is something you're likely to find on professional Digital SLR's and even some consumer grade Digital SLR's. It is generally used by sports or wildlife photographers to keep a moving subject in focus.

Autofocus Servo is normally engaged by switching focus mode to Continuous (Nikon) or "AI Servo" (Canon), half-pressing the shutter release will begin focusing and the camera will continue to focus based on it's own focus rules (and your settings) while the shutter release is half-pressed or fully depressed (actually taking shots). It's worth noting that Autofocus Servo normally also puts the camera into "release priority" mode, that is no matter what the current AF status (good lock or still searching) the camera will take a shot when the shutter release is pressed.

Here's what Nikon say about their continuous autofocus (from the D1 manual):

Continuous servo AF (release priority): The camera focuses continuously as long as the shutter-release button is held half way down, and focus will automatically track moving subjects (focus tracking). The shutter can be released even when the in-focus indicator () is not displayed in the viewfinder (release priority).

Here's what Canon say about their AI Servo autofocus (from the D30 manual):

This mode is suited for moving subjects when the focusing distance keeps changing. With its predictive AF function, the camera can also track a subject that is steadily approaching or retreating from the camera. The exposure settings are determined immediately before the picture is taken.

Batteries

The biggest bane of a digital photographer's life is batteries. The consequences of digital capture and the use of LCD's is that digital cameras are power hungry beasts. It's the price you pay for not having to reload the camera with film every 36 frames.

Most second generation (1999-2000) digital cameras take AA batteries, but if you intend on using normal off-the-shelf alkaline, you'd better think again, you'll be lucky to half fill your CompactFlash / SmartMedia card with a set of those. If your camera takes AA batteries then you're better off with a good set (or three) of NiMH rechargeables which can last between an hour and a day (depending on usage).

More recently (2000+) manufacturers have begun to supply cameras with either internally charged or externally charged Lithium-Ion or NiMH battery packs. The life of these batteries depends on their capacity, good ones providing 8Wh of power, lesser packs (often smaller and lighter) only around 4Wh.

NiCD batteries

Nickel Cadmium batteries, probably the most common and most robust rechargeables, good for on average 700 charge and discharge cycles. NiCD batteries suffer badly from memory effect (described below) which means they must be full discharged before charging and are normally only used as backup batteries.

NiMH batteries

Nickel Metal Hydride batteries, probably the most popular rechargeable batteries amongst digital camera owners, offer about 40% more capacity than NiCD plus they don't suffer from memory effect, can be charged whenever. However they're only good for about 500 charge and discharge cycles. NiMH batteries also lose about 5% of their charge a day, even when being stored.

Lithium-Ion batteries

Lithium Ion batteries, normally "the choice" of proprietary battery packs, such as those used by Sony's "InfoLithium" batteries. Lithium Ion batteries offer about twice the capacity of a similarly sized NiMH battery, however they require their own special charger and can be expensive (you can't as yet buy AA LiON batteries). They're good for about 500 charge and discharge cycles.

Memory Effect

Charging NiCD batteries before they are fully discharged will reduce the maximum capacity of subsequent charges. As the effect gets stronger when repeated often, it is called "memory effect". The accumulation of gas bubbles on battery cell plates of a battery that has only been partially discharged before recharging, a bubble reduces the plate area within the battery and thus capacity. NiCD batteries are well known for this problem.

Reconditioning

Some chargers feature a discharge / recondition button which simply discharges the battery before charging it, even more sophisticated chargers will actually negative pulse the battery to remove any built up gas bubbles.

Other cameras require rechargeable AA batteries, good brands include NexCell and GP NiMH batteries, also important is a good charger, Thomas Distributing sell the excellent, reliable and highly recommended Maha MH-C204F charger and it's bigger brother the Maha MH-C777 charger.

Buffer

A buffer inside a digital camera is RAM storage, which temporarily queues images before they're written out to the storage card to speed up the 'time between shots' and ability for the camera to shoot burst (continuous) shots. The very first digital cameras didn't have any buffer, this meant that after you to took the shot you had to wait for the image to be written out to the storage card.

Most modern digital cameras (especially those at the prosumer level) often have relatively large buffers, which allow them to operate as quickly as a compact film camera while writing data out to the storage card in the background (without interrupting your ability to shoot).

The exact location of the buffer within the cameras system affects the cameras ability to buffer images. While we can't be exactly sure (as no manufacturers have ever made this information publicly available) it's fairly easy to surmise there are two different types of buffers:

After Image Processing Buffer

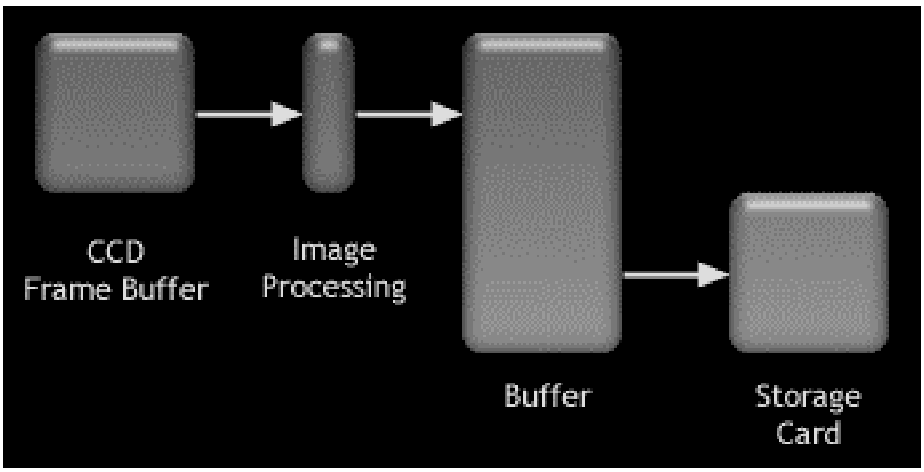

With this method the images are processed and turned into their final output format before they're placed in the buffer, this means that the number of shots which can be taken in a burst is affected by the image mode (final file size), using JPEG and/or higher compression and/or lower resolution in cameras like this will increase the number of images buffered.

Examples of cameras which use buffers like this: Canon G1, Nikon Coolpix 990, Canon EOS-D30.

Before Image Processing Buffer

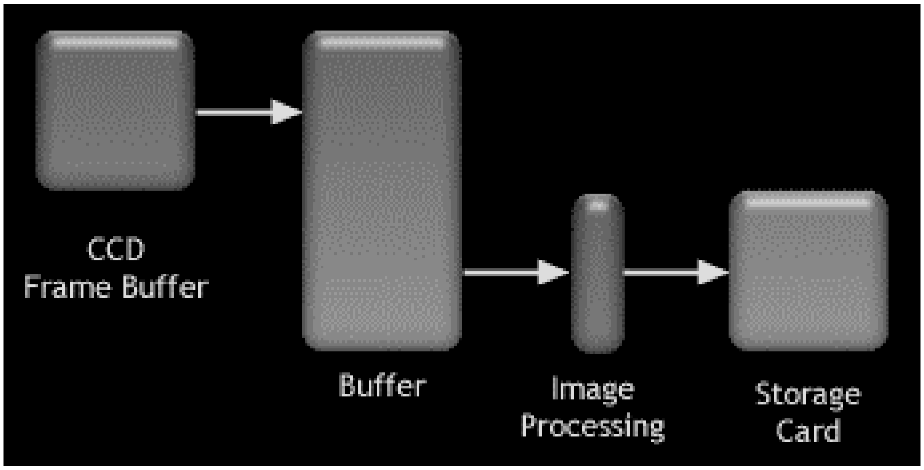

In this method no image processing is carried out and the RAW data from the CCD is placed immediately in the buffer, in parallel to other camera tasks the RAW images are processed and written out to the storage card. The problem with this method is that you're always limited to the number of frames you can shoot, no matter what image quality / size you select.

Examples of cameras which use buffers like this: Fujifilm S1 Pro, Fujifilm 4900Z, Olympus C-3030Z.

Burst (Continuous)

Burst or often named Continuous mode shooting can be likened to a film SLR with a motorwind, the digital cameras ability to take several shots immediately one after another. The speed and number of shots differs greatly from one camera to another, though most prosumer digital cameras can manage about 3 fps for around 6 frames at full resolution and more at lower resolutions.

The cameras ability to shoot quickly is a function of it's shutter release and image processing systems, the number of frames that can be taken is defined by the size of the internal memory buffer where images are stored before they're written out to the storage card (as storage card devices are currently too slow to take the high data rates created by burst shooting). Some cameras can data out of their buffer at the same time as new frames are being stored in it, thus you're never actually stopped from shooting but the frame rate slows down (as you're waiting for buffer space to free).

For really serious burst ability you need to look at the professional Digital SLR's such as Nikon's D1 (up to 4.5fps for 21 frames) or the Kodak DCS series (approx. 3.5fps for 12 frames).

Color Filter Array

Photodiodes, the tiny light sensitive materials used to measure the amount of light for each pixel on a CCD are essentially monochrome devices, that is they can't themselves tell the difference between different wavelengths of light. To produce a color image a CFA (Color Filter Array) is placed over the monochrome sensor pixels, in reality the CFA is made up of very thin layers of colored dye. This CFA filters out all but the chosen color for that pixel. Software interpolation later produces a full color for a pixel based on the value of surrounding pixels.

The GRGB Bayer Pattern is the most common CFA used, it is used in most consumer digital cameras (except for Canon up until the S100 and Nikon on the Coolpix series that use a CYGM CFA). A primary color (GRGB) Bayer Pattern is produced by placing two layers of colored dyes over each other as such:

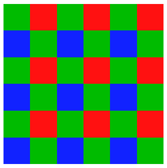

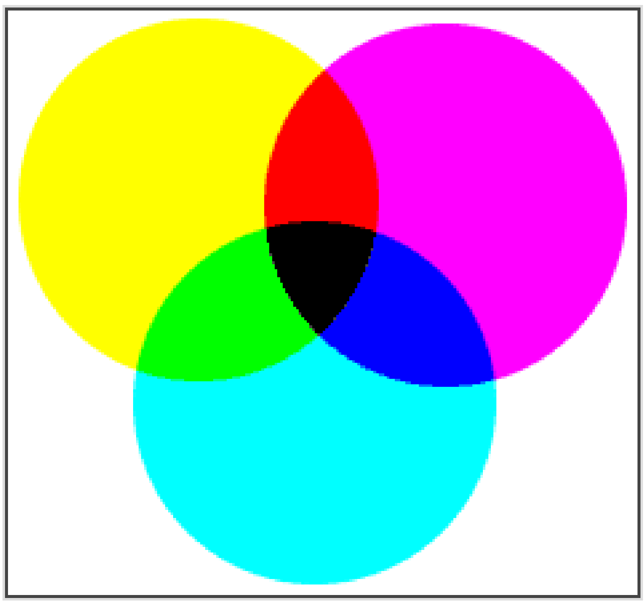

Red = Yellow + Magenta
Green = Yellow + Cyan
Blue = Magenta + Cyan

On a 2048 x 1536 pixel CCD using a GRGB CFA the following is true:

* 1024 x 768 pixels are RED (768,432)
* 1024 x 768 pixels are BLUE (768,432)
* 1024 x 1536 pixels are GREEN (1,572,864)

On a 2048 x 1536 pixel CCD using the CYGM CFA the following is true:

* 1024 x 768 pixels are CYAN (768,432)
* 1024 x 768 pixels are GREEN (768,432)
* 1024 x 768 pixels are MAGENTA (768,432)
* 1024 x 768 pixels are YELLOW (768,432)

How the camera creates the final image

So the output from the CCD is made up of groups of color intensities, the camera's internal processing algorithms then combine each of these colored pixels with the value of their neighboring pixels to produce the final image:

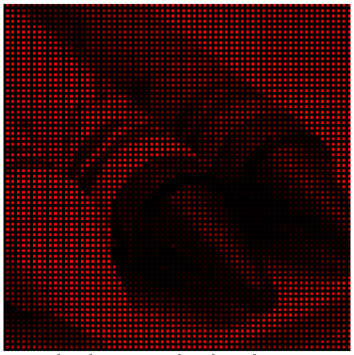

Red channel pixels

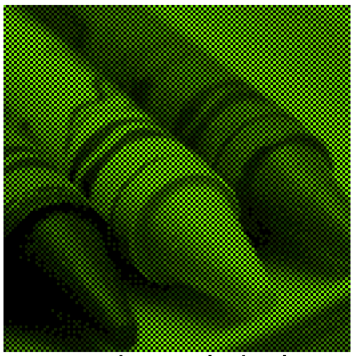

Green channel pixels

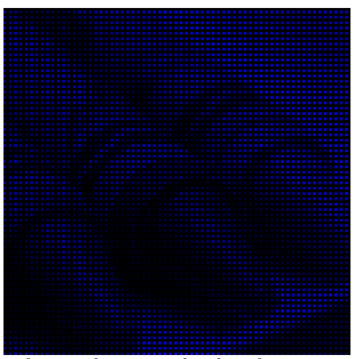

Blue channel pixels

Combined

As you can see, the combined image isn't quite what we'd expect but we could definitely make out and image and the colors of each of the items in the scene. If you stand away from your monitor your eyes will combine the individual red, green and blue intensities to produce a (dim) color image. Finally, the combined image is run through a demosaicing algorithm, which combines the color values of a pixel, and it's eight neighbors to create a full 24-bit color value for that pixel:

Connectivity

A digital camera's connectivity defines how it can be connected to other devices (computers primarily) for either the transfer of images or remote control of the camera.

Image Transfer

Early digital cameras used RS232 (serial) connections for image transfers, in my honest opinion this was probably the single worse thing about owning a digital camera in late 1998 - 1999. Most consumer grade digital cameras now feature USB connectivity; this provides a relatively fast transfer rate (up to 500Kbyte/s) to a wide range of computer systems on the PC and Mac platforms. Manufacturers generally bundle such cameras with cables and driver software. One step up from USB is Firewire (IEEE1394) seen mostly on professional digital cameras, this offers considerably faster transfer rates but requires that you computer has a Firewire connector or you're willing to fit it out with Firewire. These are the kinds of transfer rates you can expect from different devices:

Device	**Approx. Transfer Rate**
Digital Camera USB	~ 350 KB/s
Digital Camera Firewire	~ 500 KB/s
USB Card Reader	~ 500 KB/s
Lexar Media Jumpshot CF->USB	~ 780 KB/s
SCSI Card Reader	~ 1,000 KB/s
PCMCIA Card Adapter (laptop/notebook)	~ 1,300 KB/s
Firewire Card Reader	~ 2,200 KB/s

Remote Control

More recently manufacturers have started to bundle cameras with remote control software (by recently I mean Fujifilm's S1 Pro, Canon EOS-D30). These remote control applications use the same USB connection for image transfer to control the camera, normally with time lapse and timed capture features and even allow images to be taken and stored directly onto your local hard disk.

Video Output

Additionally most digital cameras also provide video (and sometimes audio) output for connection to a TV or VCR. More flexible cameras allow you to switch output between the PAL and NTSC video standards, cameras with InfraRed remote controls make it easy to do slideshows for friends and family from the comfort of your armchair.

Effective Pixels

Most digital camera manufacturers use the total number of pixels of the sensor used as an indication of the camera's resolution. However, the true pixel count is the number of recorded pixels (not strictly effective pixels but now a common term and used here to represent this number).

A sensor has four measurements of its pixel count (figures below are from Sony's ICX252AQ 3.34 megapixel CCD used in a large number of 3 megapixel cameras):

Total number of pixels	2140 x 1560 (3.34M)
Number of read pixels	2088 x 1550 (3.24M)
Number of active pixels	2080 x 1542 (3.21M)
Recommended recorded pixels	2048 x 1536 (3.14M)

As you can see the camera doesn't record an image using the total number of pixels, not even the total number of active pixels but a "crop" of that. Why? Well, some of the vertical columns and horizontal rows on the edges of a sensor are covered by a black dye (known as Video Signal Shading), for a digital camera to know "what black is" it needs to take a "dark current" reading, this is the base value (not necessarily zero) of a pixel which is completely black.

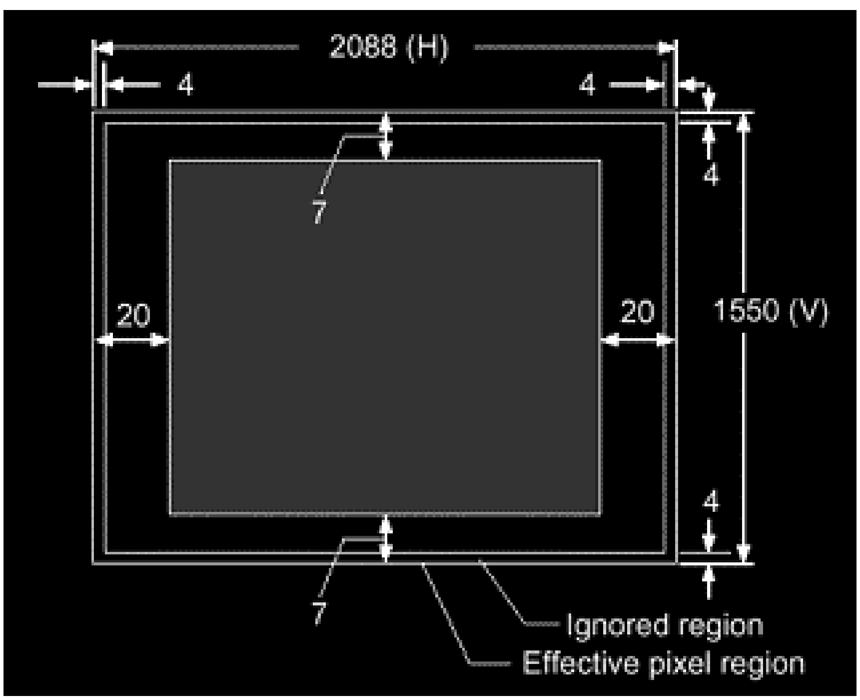

The diagram above shows the number of read pixels (2088 x 1550) and the 8 rows 8 columns of masked region around this reducing the amount of pixels actually exposed to 2080 x 1542 (the active pixels). Within this manufacturers have the flexibility to choose the output image size, most will of course choose a common resolution which fits into a standard aspect ratio (such as 2048 x 1536 for 3 megapixel 4:3 aspect ratio).

We list both the CCD pixel count and the effective pixel count in our specifications database, it is the effective pixel count which matters. A good example of this is Sony's DSC-F505V which uses the same CCD described above and has "3.34 megapixels" silk-screened on the side but only uses 2.6 megapixels (1856 x 1392) of these for the final output image, why? Because they fitted this new CCD into last year’s body and as the CCD was slightly larger the lens wasn't able to cover the whole CCD frame.

LCD

The LCD of a digital camera is the screen used to review, preview and even act as a large size viewfinder. Many consumer digital cameras allow you to use the LCD to frame the scene providing a live video feed of the image to be captured. The LCD screen is also used to review images taken and change camera settings.

LCD screens come in many different forms, but on digital cameras are normally between 1.5 and 2.0" diagonally using TFT technology. The better screens have an anti-reflective coating and/or a reflective sheet behind the LCD to allow the LCD to be used in bright outdoor daylight (example: Sony DSC-F505, F505V). Some viewfinders can be flipped out of the body or angled up or down to make it easier to take low angle or high angle shots.

LCD Viewfinders

Some newer digital cameras now feature LCD viewfinders, these viewfinders have a very small LCD screen (often just 0.5" diagonally) to project the live view image captured by the CCD. These viewfinders are becoming more popular but often lack the accuracy and resolution of a true TTL viewfinder.

Sensor (CCD/CMOS)

The sensor is the chip (CCD or CMOS) which records light falling on it, in the digital camera. It is the device which actually captures "the picture" (a digital camera's film). Originally developed for video applications they have progressed in resolution and color accuracy to a stage where we now find multi-megapixel cameras common.

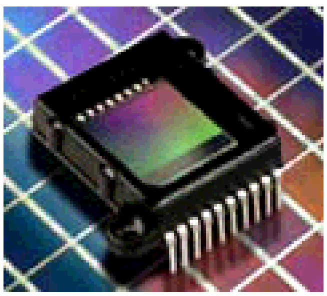

A common analogy is an array of buckets on conveyor belts, the raindrops represent the photons of light falling onto the CCD surface and being captured in 'bit buckets' (photosites; pixels). The conveyor belts which empty them are known as the shift registers, in progressive CCDs, the CCD is read: one horizontal line, shift the vertical down one pixel, one horizontal line, and repeat...

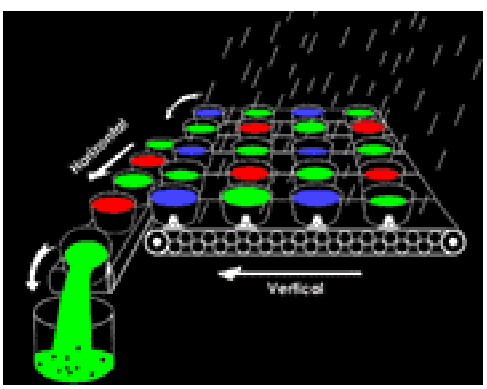

Even though a chip will profess to being 3 megapixels, it is not true that each pixel can in fact capture any color, the CCD itself is a monochrome device and therefore requires a color filter (usually produced by painting directly onto the surface of the CCD using dye) to produce a colored image. See Color Filter Array.

The dynamic range of a sensor is defined by the depth of the wells used in the photosites; this can be imagined as the depth of the bucket, the number of electrons it can hold before overflowing. Noticeably professional digital SLRs and even modern consumer grade digital SLRs have considerably more dynamic range than consumer digital cameras. See Dynamic Range, AD Converter.

Interline Transfer sensor

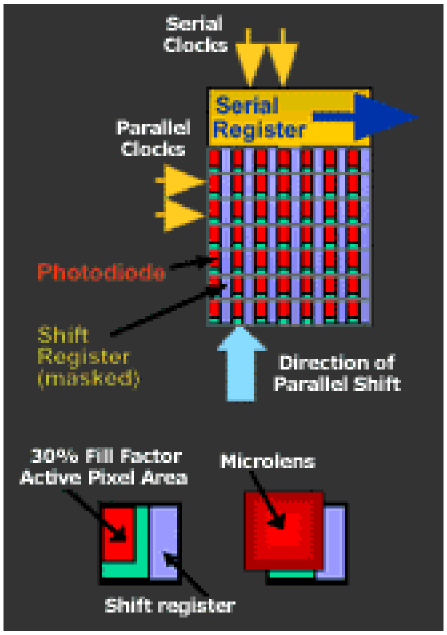

Your typical consumer digital camera uses what is called an Interline Transfer CCD, put simply the CCD can itself control the start / stop of when it measures light falling on it, otherwise known as an electronic shutter, it does so by shifting values out of the photodiodes into "shift registers" then pushing all of that data out as a final image. The advantages to Interline Transfer CCDs are simply that, they can be controlled by software and don't require a mechanical shutter (though are often used in conjunction with one) and can produce a video feed output (a requirement for a live preview LCD feed).

Because of the extra electronics required around each pixel the "fill factor" (size of the photodiode) tends to be quite small (about 30% of the pixel area). To get around this Interline Transfer CCD manufacturers place a layer of "microlenses" (click here for electron microscope image of microlenses) over the CCD to capture more light and focus it onto the smaller photodiode area which then gives them an better effective fill factor of about 70%.

Full Frame sensor

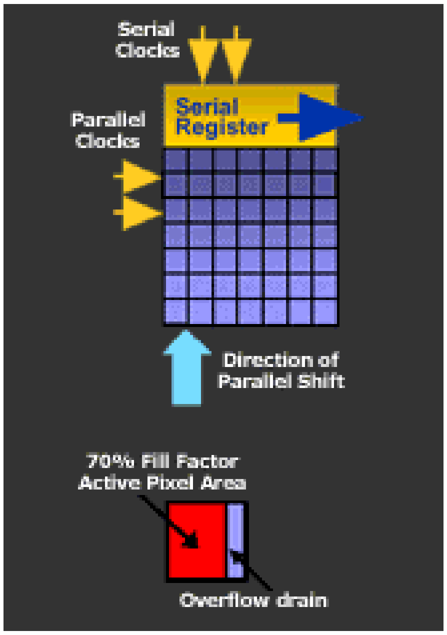

CCDs used in professional cameras tend to be Full Frame Transfer, they don't have a shift register, this means that a mechanical shutter is absolutely required to control the start / stop measurement of light. The shutter is opened and then closed again (say 1/60s later), the whole CCD shifts data off itself into the serial register where it's processed as the "RAW" image. As Full Frame CCDs are simpler (don't have shift registers and associated electronics around each photodiode) they have a much better Fill Factor (around 70%) and don't require or use microlenses.

The disadvantage is that you can't get a video feed out of them which is the main reason we don't see more manufacturers using Full Frame CCDs (we're all too used to our LCD preview).

Pro's & Con's associated with each CCD type:

Full Frame CCD			**Interline Transfer CCD**		
+	High image quality		+	Good image quality	
+	High sensitivity		+	Good sensitivity when using microlenses	
+	High dynamic range		+	Low noise	
+	Larger sizes		+	High frame rates / electronic shutter	
+	No microlenses		+	Video feed capable	
-	Not capable of video feed		+	Don't need mechanical shutter	
-	Top shutter speeds limited by mechanical shutter		-	Microlenses can cause aberrations	
-	Require mechanical shutter				

Storage Card

The storage card in a digital camera is it's film, it's the removable storage device which holds images taken with the camera (a few low-end digital cameras don't have removable storage cards but instead have a built-in flash RAM unit). There are currently several competing storage card formats used in today's digital cameras:

PCMCIA PC Card

Storage cards of this type have exactly the same dimensions and electrical contacts as PCMCIA cards in notebooks, and can be used directly in a PCMCIA slot without any adapter. These cards are only found on high-end professional digital SLRs, because of their large size they're not compact enough to be used in a "compact" digital camera. Available in three different physical sizes: Type I, II and III, both flash memory and hard drives as a large range of capacities up to several GB.

Dimensions: 85.6 x 54.0 x (Type I: 3.3mm, Type II: 5.0mm, Type III: 10.5mm)

Compact Flash Type I

Probably the most common format used today, found in both digital cameras and other devices (PDAs, network applications etc.), small format and lightweight they have the same electrical contacts as a PCMCIA card and can be used in a PCMCIA slot with a simple adapter. They are available in capacities up to 256 MB.

Dimensions: 43.0 x 36.0 x 3.3mm

Compact Flash Type II

The only difference between Type I and Type II is the size of the package; slightly thicker at 5.5mm compared to Type I's 3.3mm. The large package size allows for larger capacities, notably IBM's successful Microdrive is a Type II device with capacities of 340, 512 and 1 GB. As with Type I devices Type II devices can be use directly in a PCMCIA slot with an adapter.

Dimensions: 43.0 x 36.0 x 5.5mm

SmartMedia

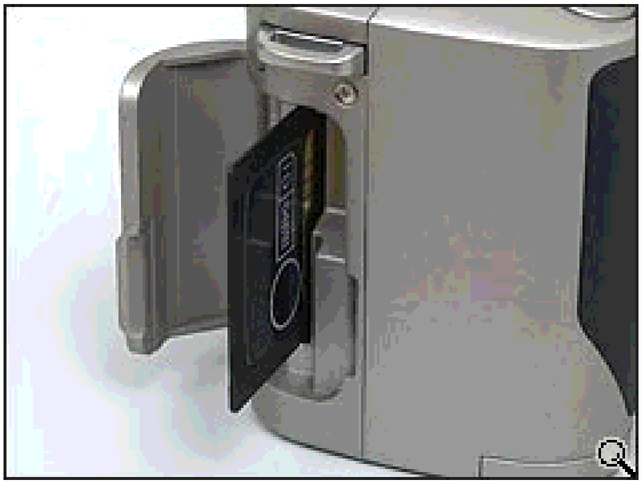

SmartMedia is considerably thinner than Compact Flash, without the controller chips found in a Compact Flash card all that's found inside the SmartMedia's package are the flash memory chips themselves. Unfortunately this means that the controller chip in the digital camera must be programmed to recognize the storage card (a problem when 64MB SmartMedia cards were first introduced, and likely to be a problem when the new 128MB SmartMedia cards are introduced).

Dimensions: 45.0 x 37.0 x 0.76mm

Sony MemoryStick

When Sony hit the market with their chewing gum sized MemoryStick, nobody thought it would be successful until Sony sold the rights to manufacture MemoryStick and MemoryStick compatible devices to other companies. Now Sony have done a good job of ensuring there are plenty of devices already available which use the MemoryStick. They are available in capacities up to 128 MB with 256 MB on the horizon. Only currently found in Sony Digital Cameras.

Dimensions: 50.0 x 21.5 x 2.8mm

Other Types

There are a few other types of storage card which haven't yet made their way into commercially available digital cameras, we'll update this section when they do.

Digital Imaging

Artifacts

Artifacts simply refer to distortions of an image produced by either the image sensor (CCD/CMOS) and/or optical system, internal image processing algorithms or compression algorithm (JPEG).

Image Sensor / Optical system artifacts

It's often difficult to attribute artifact effects to one part of the optical system of a digital camera, with most consumer cameras the whole optical pathway is sealed and thus technical analysis of what's going on here is something only the manufacturer's R&D departments can answer. On higher end digital SLRs with large pixels an anti-alias filter must be used to reduce the effects of higher frequency detail creating moiré patterns. However, these anti-alias filters often introduce unwelcome artifacts into the image. Also see noise, chromatic aberrations, blooming.

Internal image processing algorithm artifacts (& sharpening)

Often caused by mathematical errors and a combination of optical system artifacts, the internal processing algorithms can amplify these small errors to something much more visible. The other type of image processing artifact stems from over sharpening of the image internally, normally visible as a black or white halo around contrasting areas of the image (this can normally be avoided by turning down the sharpening settings on the camera, if available).

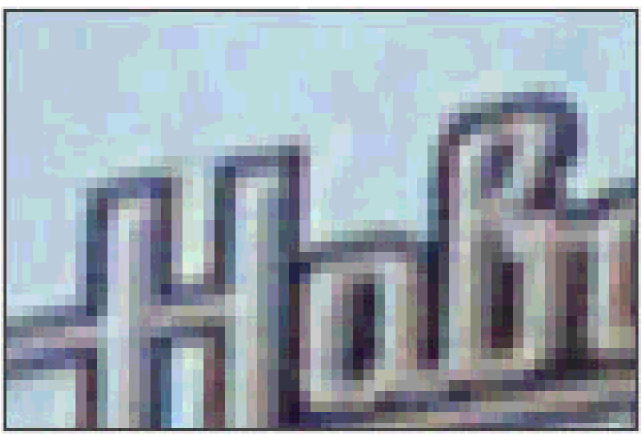

Example above shows visible sharpening artifacts as a white halo around the top edge of this black lettering.

JPEG compression artifacts

When you save images in JPEG format the algorithm "loses" (AKA lossy compression) certain information to produce a smaller file size. The more "heavily" you compress (smaller file) the more information you lose and the more artifacts are created in the image. JPEG is particularly susceptible to artifacts because of the way it attempts to maintain details (edges) against large plain color areas. Because JPEG analyses the image in 8x8 blocks, these artifacts can sometimes appear with sharp "square" edges. JPEG also doesn't handle very noisy images very well, because of the amount of information in a very noisy image JPEG has to throw away a lot of information which introduces yet more artifacts and can make the whole situation a whole lot worse.

Below is a rather severe example of JPEG artifacts, hardly any digital cameras used compression as aggressive as this but this example should serve to demonstrate why it's always important to use as low a compression ratio as you can. Images below cropped and magnified 200%.

When editing an image in several sessions, it is recommended to save the intermediate image in an uncompressed format (tiff, bmp, and jpeg maximum quality) to avoid cumulation of JPEG losses. If you save for instance an image at Photoshop JPEG Quality 4, close it, open it again and save it again at Photoshop JPEG Quality 4, the file size will not reduce further, but quality will have degraded further. So only compress after all editing is done.

Blooming

Each pixel (photosite) on a digital camera sensor (CCD/CMOS) has a limit as to how much charge it can store. Blooming (or Streaking) is the name given to an overflow of charge from an oversaturated pixel (photosite) to the next on the sensor. This problem is addressed with the addition of "anti-blooming gates" which can be thought of as vertical drain ditches running beside each row of pixels, these gates allow the overflowing charge to run away without affecting surrounding pixels.

Anti-blooming gates, while mostly successful (and certainly for more modern sensors) blooming can still be a problem in very extreme exposures (very bright edge against a virtually black edge) and is typically visible as either a vertical streak or white halo extending for several pixels. The effects of blooming often amplify the visibility of chromatic aberrations. Blooming is really good at destroying the detail of leaves shot against a bright sky.

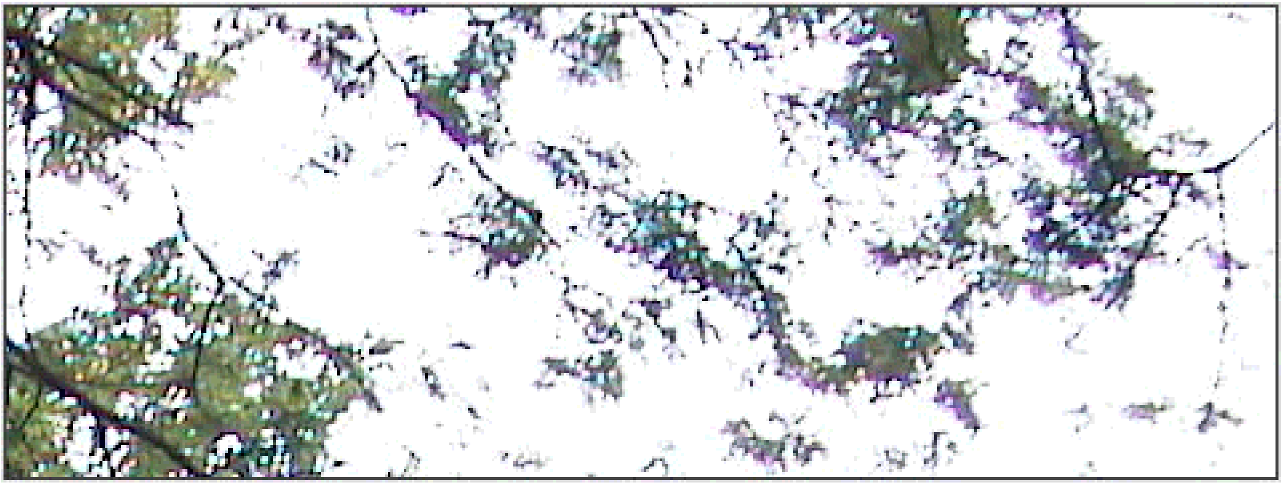

Color Space

Color spaces describe how the red, green and blue primaries are mixed to form a given hue in the color spectrum. Since it is not possible to represent every color in the visible spectrum *exactly* by mixing amounts of red, green and blue, color spaces allow us to change how we define red, green and blue (and white) to get better color reproduction. By "tweaking" primaries in this way, we can maximize how many colors *can* be accurately represented on monitors, printers, etc. by matching the color space closely to what the device is capable of reproducing. There are also color spaces designed for photo editing that don't match *any* particular device, but are instead designed to give more "coverage" of the overall color gamut (spectrum of visible light) for the purpose of photo editing.

Dynamic Range

Dynamic range is the ratio between the brightest and darkest recordable parts of an image or scene. A scene that ranges from bright sunlight to deep shadows is said to have a high dynamic range, while indoor scenes with less contrast will have a low dynamic range. Note that depending on the scene contrast, it may or may not be possible to capture the entire range with a digital camera. In recording scenes with very high dynamic range, digital cameras will make compromises that allow the capture of only the part of the scene that is most important. This compromise is needed because no camera or output device of any kind (including the human eye) can reproduce the nearly infinite dynamic range that exists in real life.

Example of an image with good dynamic range, all bright sky and wall detail is visible without destroying shadow detail, a 'wide dynamic range'.

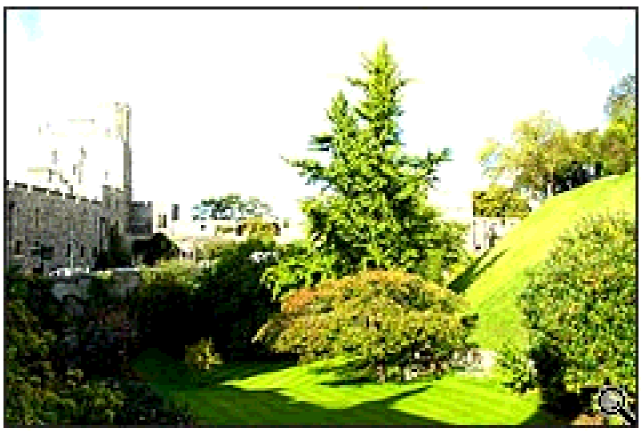

Simulated example of an image with poor dynamic range, sky and wall details are blown out because the camera metered for shadow and didn't have the dynamic range to maintain the sky.

Gamma

A monitor (CRT) does not respond in a linear fashion such that the luminance level corresponds directly to input (brightness values in an image). Instead of responding in a 1:1 ratio of input to monitor brightness level, a monitor normally responds on a nonlinear curve due to hardware limitations. The response of a typical PC monitor can be characterized by taking the input luminance level (a factor between 0 and 1.0) and raising this value to a power of 2.2 to 2.5. The monitor response is therefore generally darker than intended, especially at the dark end of the scale.

Gamma Correction

Gamma correction is a method by which images are modified to correct for nonlinear response of a particular hardware device (see gamma above). Gamma correction is designed to compensate for the gamma response of a device by raising the input luminance levels (0 being the darkest and 1.0 being pure white) to the power of 1/2.2 to 1/2.5 depending on the gamma of the device. Raising luminance levels to the power of 1/gamma will properly account for the nonlinear response of a device by converting luminance levels so that a linear response is achieved; brightening the image to appear as it did in the original scene.

Histogram

A histogram function on a digital camera allows the photographer to quickly and easily see the accuracy and "spread" of the exposure over the cameras grayscale range (dynamic range). The histogram itself is a graph of brightness along the horizontal axis (black to white) and the number of pixels at each brightness level on the vertical axis. Most digital SLRs and some newer prosumer digital cameras feature a histogram review mode, understanding what the histogram is telling you about your exposure is important to get the most out of your digital camera. Examples of histograms below.

Implementing a histogram as part of the cameras review mode provides invaluable information. The examples below were of real life images; in this case the histograms were extracted using Adobe Photoshop.

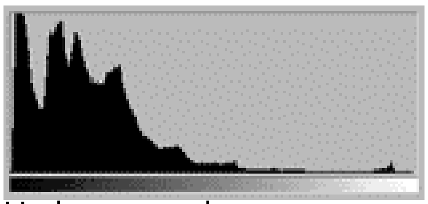
Underexposed

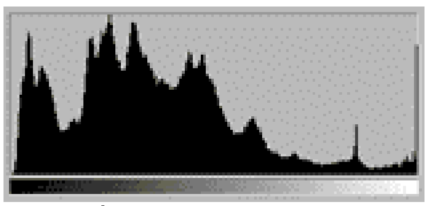
Normal exposure

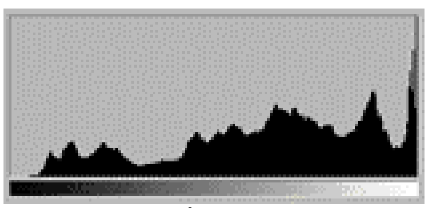
Overexposed

If the "weight of pixels" is predominantly in either the bottom or top part of the grayscale then the image is likely to be under or over exposed. Of course the histogram is only another tool in assisting you to create the shot you want.

Noise

Noise can be summarized as the visible effects of an electronic error (or interference) in the final image from a digital camera. Noise is a function of how well the sensor (CCD/CMOS) and digital signal processing systems inside the digital camera are prone to and can cope with or remove these errors (or interference).

Visible noise in a digital image is often affected by temperature (high worse, low better) and ISO sensitivity (high worse, low better). Some cameras exhibit almost no noise and some a lot and all the time. It's certainly been the challenge of digital camera developers to reduce noise and produce a "cleaner" image, and indeed some recent digital cameras are improving this situation greatly, allowing for higher and higher ISO's to be used without too much noise.

Noise can also affect certain color channels more than others, this is because a typical digital camera sensor (CCD/CMOS) is more sensitive to certain primary colors than others (often sensors are less sensitive to blue light) and so to compensate, these channels are amplified more than the others. Noise is also often amplified by the JPEG compression algorithm, which reacts badly to a very noisy image (often introducing hue errors, which weren't in the original noisy image). See also artifacts.

There are several techniques for cleaning a noisy image, and several products that have been developed to specifically perform this task. Check out Quantum Mechanics by Camera Bits and the noise filter in QImage Pro by Michael Chaney. There's a technique here (PDF) from Adobe for noise removal using Photoshop.

The crops below are taken from my Canon EOS-D30 review, they show that there's almost no visible noise at ISO 100, but at ISO 1600 there's a fair amount of visible noise in the image.

ISO 100 image

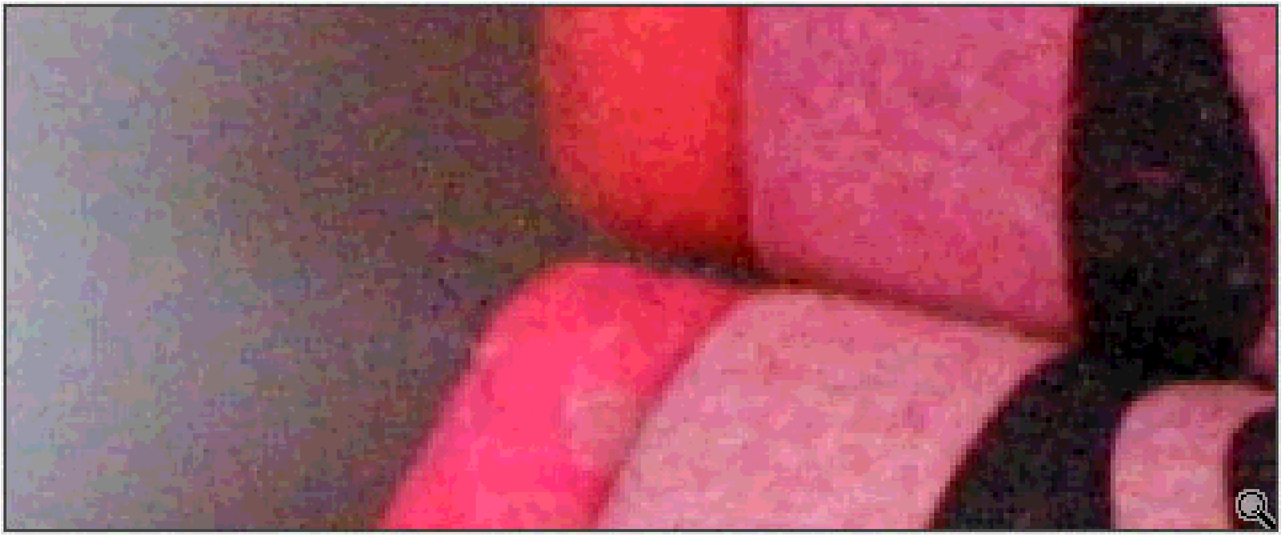

ISO 1600 image

Long Exposure "Stuck Pixels"

Another type of noise often referred to as "stuck pixels" or "hot pixels" occurs only with long exposures (1-2 seconds or more) and appear as fixed colored dots (slightly larger than a single pixel). These stuck pixels can be fairly successfully removed by taking a "dark frame" either before or after the main shot and subtracting this from the original shot, this technique is detailed in noise reduction.

Noise Reduction

Noise reduction is a system for removing unwanted noise from a digital image. It falls into two main categories, reduction or removal of noise from high ISO images and reduction or removal of noise from long exposure images (with "stuck pixels").

High ISO noise reduction

So far we have only seen this feature implemented in the Kodak DCS620x, this method of noise reduction relies on a sophisticated algorithm and on the processing power of the acquiring computer, noise reduction is carried out at the time of conversion of the RAW image for importation. See the Kodak DCS620x review for more details.

Long exposure ("stuck pixel") noise reduction

This method of noise reduction relies on a second frame known as the "dark frame" being taken with the shutter closed for the same amount of time as the main image, this dark frame will be taken either before or after the main image and will be used to identify and subtract "stuck pixels". Kodak has implemented this in their consumer grade digital cameras for some time, and more recently in their new DC4800 digital camera. Canon also implemented this method of noise reduction in the Canon EOS-D30 digital SLR.

This method of noise reduction can also be implemented in a software package such as Photoshop. Simply take the original image, then a second image with the lens covered with exactly the same exposure (this becomes the "dark frame"). Open the original image in Photoshop, paste the "dark frame" as a new layer, give it a slight Gaussian Blur (around 0.3 pixels) then change the layer options to "Difference". This will make a fairly good job of removing stuck pixels from the noisy original image. Note: this isn't perfect and will still leave some black "pits" in the image. The examples below were taken with a Nikon Coolpix 990 and the exposure was 8 seconds at F2.9.

Original Image Dark Frame Cleaned Image

Jaggies

Hardly a technical term, jaggies refer to the visible "steps" of a diagonal line / boundary in a digital image. These steps are simply a consequence of the regular, square layout nature of a pixel. To reduce the visibility of jaggies, lines are aliased against the background or along the boundary so as to look like a smoother line. Jaggies are normally more obvious on lines at slight angles (such as wires between telegraph poles or rooftops against the sky).

Jaggies are (unfortunately) made more visible by sharpening the image (as sharpening is often based on selective contrast adjustment that removes aliasing). Digital Cameras often employ fairly strong sharpening as a standard part of their image processing algorithms, even with digital cameras which allow you to control the sharpening levels their lowest sharpening is still having an effect on the image. The image below shows five squares, each one at a slightly different angle, the first not rotated at all the rest rotated by 5, 10, 30 and 45 degrees. Without blowing up this image it looks perfectly fine with smooth clean edges.

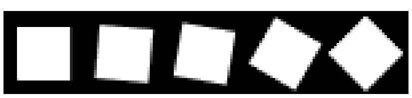

Magnified 200% it's now fairly easy to see how the aliasing works and how it smoothes the edges of lines (digital cameras don't actually perform aliasing but it's a side effect of the interpolation techniques used to reconstruct the image from the RAW sensor data).

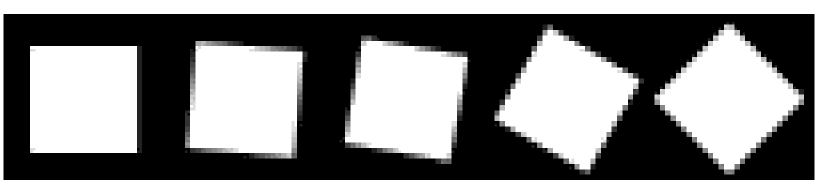

JPEG

The JPEG image format (actually JPEG is strictly speaking the name of the compression algorithm) is the predominant format used by digital cameras. Even professional cameras will have a JPEG mode and despite its lossy compression it has become the de facto image format for electronic storage of photographic images. The reason for this is this compression algorithm's ability to reduce an image file size by 8 or 10 to 1 without any degradation in image quality (to the human eye at normal viewing magnification). A 1.5 MB 3 megapixel JPEG image will be hard to distinguish from the same image saved as a lossless TIFF weighing in at 9.2 MB.

I'll not delve too deeply into exactly how JPEG works (there are too many papers available on the web which cover that in far more detail) but I'll summarize roughly what happens:

1. The image data is converted from RGB (if it was already in the RGB color space) to YUV; YUV is a color space scheme that stores information about an image's luminance (brightness) and chrominance (hue) separately. Since the human eye is much more sensitive to luminance than chrominance, you can afford to discard much more information about an image's chrominance, especially the higher frequencies.

2. Image is broken up into 8x8 blocks that are then processed independently.

3. Each block is processed by the DCT (Discrete Cosine Transform) algorithm "The ultimate goal of the DCT is to represent the image data in a different domain using the cosine function. The image data is transformed into numerous curves of different sizes. When these curves are put back together, through inverse Discrete Cosine Transform, the original image (or an extremely close approximation to it) is restored."

4. Quantization - two tables of quantization constants are created, one for luminance and one for chrominance. These constants are calculated based on the image quality variable (the "JPEG quality").

5. Encoding - elements of representing low frequencies are moved to the beginning of the matrix and elements representing high frequencies (more likely to be zeros and thus better to compress together) are moved towards the end of the matrix. The matrix is then compressed (generally) using Huffman encoding and written to the output file.

See also artifacts.

When editing an image in several sessions, it is recommended to save the intermediate image in an uncompressed format (tiff, bmp, and jpeg maximum quality) to avoid cumulation of JPEG losses. If you save for instance an image at Photoshop JPEG Quality 4, close it, open it again and save it again at Photoshop JPEG Quality 4, the file size will not reduce further, but quality will have degraded further. So only compress after all editing is done.

References

JPEG Algorithm and Associated Data Structures, Mark D. Schroeder, University of North Dakota. http://people.aero.und.edu/~mschroed/jpeg.html

JPEG FAQ http://www.faqs.org/faqs/jpeg-faq/part1/

TIFF

TIFF is actually an acronym for Tagged Image File Format and is an extremely complex and flexible image format. In digital cameras TIFF has been used to provide a "lossless" image format, that is one which doesn't throw away information in the compression process.

TIFF for photographic images usually uses an 8 or 16-bit per color channel storage method, TIFF also supports lossless compression (LZW compression), unfortunately this is patented by Unisys so many digital camera manufacturers don't use it and instead output uncompressed images.

TIFF isn't a very good choice for photographic images stored on a digital camera storage card, primarily because it creates very large files, for example a single TIFF image file from a 3 megapixel digital camera will be in excess of 9 MB. A far better image format (but not widely accepted, nor standardized) is RAW.

Note: Despite the filename extensions Kodak's Professional DCS digital SLR's TIFF files are actually a proprietary Kodak RAW image format.

RAW Image Format

The most common image format amongst digital cameras is JPEG; it's a format that produces relatively small files from large amounts of image data by discarding certain information, as JPEG uses a "lossy compression algorithm". The only other common alternative is TIFF, this produces an uncompressed 24-bit per pixel image often in the multiple megabytes, certainly for a 3 megapixel camera in excess of 8 MB per image, not really practical. A little background: each pixel of a CCD can only see one color, depending on the CFA (color filter array) placed over the CCD this is either Red/Green/Blue or Cyan/Magenta/Green/Yellow. The cameras internal image-processing engine then interpolates colors from the value of neighboring pixels to calculate a full 24-bit color for each pixel.

RAW is simply the raw data as it comes directly off the CCD; no in-camera processing is performed. Typically this data is 8, 10 or 12 bits per pixel. The advantage being that file sizes are considerably smaller (e.g. 2160 x 1440 x 12 bits = 37,324,800 bits = 4,665,600 bytes), the image has not been processed or white balanced which means you can correct the image, and it's a better representation of the "digital negative" captured. The disadvantage is you can't open these image files with a normal photo package without using an "acquire module" (a plugin, typically TWAIN, which can open / process such images).

RAW image format has actually been around for quite a while, Canon had a RAW format back in the old PowerShot range and more notably on the Pro 70, all of Kodak's DCS Pro series shoot in a proprietary RAW format (despite the TIFF extension), Nikon's D1 also has a RAW format. Canon has resurrected RAW format for the EOS-D30 and G1.

Advantages of RAW format

* A true "digital negative", untouched by cameras processing algorithms
* No sharpening applied
* No gamma or level correction applied
* No white balance applied
* No color correction applied
* Lossless yet considerably smaller than TIFF
* Records data over a wider bit range (typically 10 or 12 bits) than JPEG or 8-bit TIFF

Disadvantages of RAW format

* Requires proprietary acquire module (typically TWAIN) or plugin to open images
* Images can take 20-40 seconds to process on an average machine
* No universally accepted RAW standard format, each manufacturer (even each camera) differs

Resolution

And we're not talking about the advertised resolution! If you've not yet read about sensors and color filter arrays it's advisable to do so before understanding resolution fully. When we talk about resolution in our reviews, we're not talking about the number of pixels on the sensor (although more is normally better), and we're not talking about the output image size (the sensors effective pixel count always gives the optimum resolution).

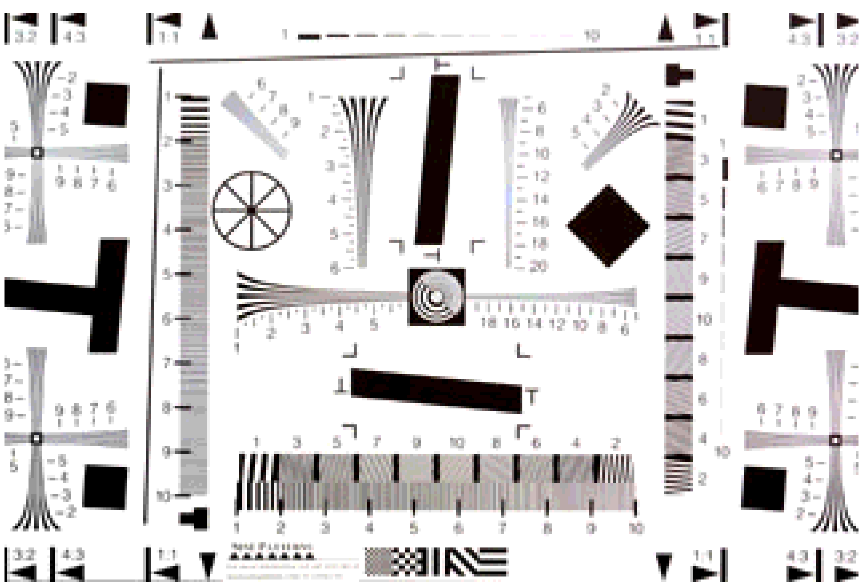

We measure resolution using the now widely accepted PIMA/ISO 12233 camera resolution test chart (specifications available here in Excel format). This chart is excellent not only for measuring pure horizontal and vertical resolution but also the performance of the sensor with frequencies at various angles; it also offers a good reference point for comparison of resolution between cameras.

We started using this chart at the beginning of this year and now make the chart available for every camera that comes through our test labs, both in the camera review and our extensive camera database.

Resolution from this chart is always measured in lines per picture height (to keep the pixels square), the numbers seen on the chart refer to hundreds of lines, so the label "12" refers to 1,200 lines per picture height.

Most recent 3 megapixel digital cameras are capable of resolving about 1,100 lines per picture height, with a 4:3 aspect ratio that would give us 1,466 x 1,100 or about 70% of the width / height pixel count of the sensor. Why only 70%? Because of the interpolation of the data required from the sensor and limitations of optics used / required to create an incredibly sharp image on such a small sensor area.

Interpolation

Interpolation (sometimes called resampling) is an imaging method for increasing the size of a digital image. Some digital cameras use interpolation to produce a larger image than the camera's sensor captured or to create digital zoom, most photo packages also support some method of interpolation for resizing images. How smoothly images are enlarged without introducing jaggies (pixelation) depends on the sophistication of the algorithm.

The examples below are all 450% increases in size of this 106 x 40 crop from an image.

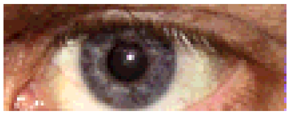

Nearest neighbor interpolation

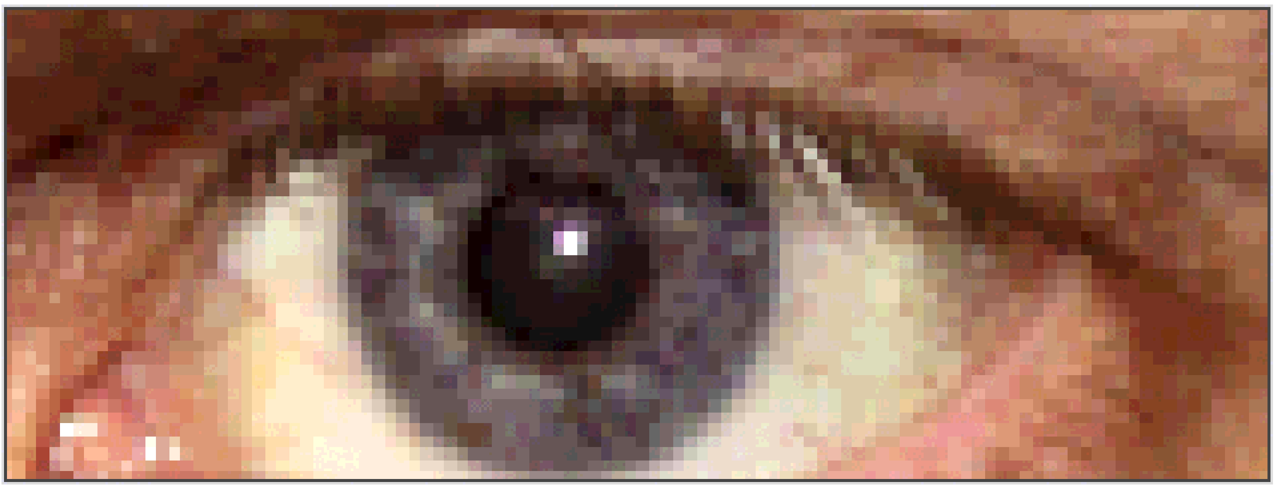

When increasing (or decreasing) the size of a digital image the simplest way to do it is called "nearest neighbor" interpolation, this simply says take the color of the new pixels nearest neighbor and use it for this new pixel, this results in pixelization and visible jaggies (but is often useful for closely examining a digital image).

Bilinear interpolation

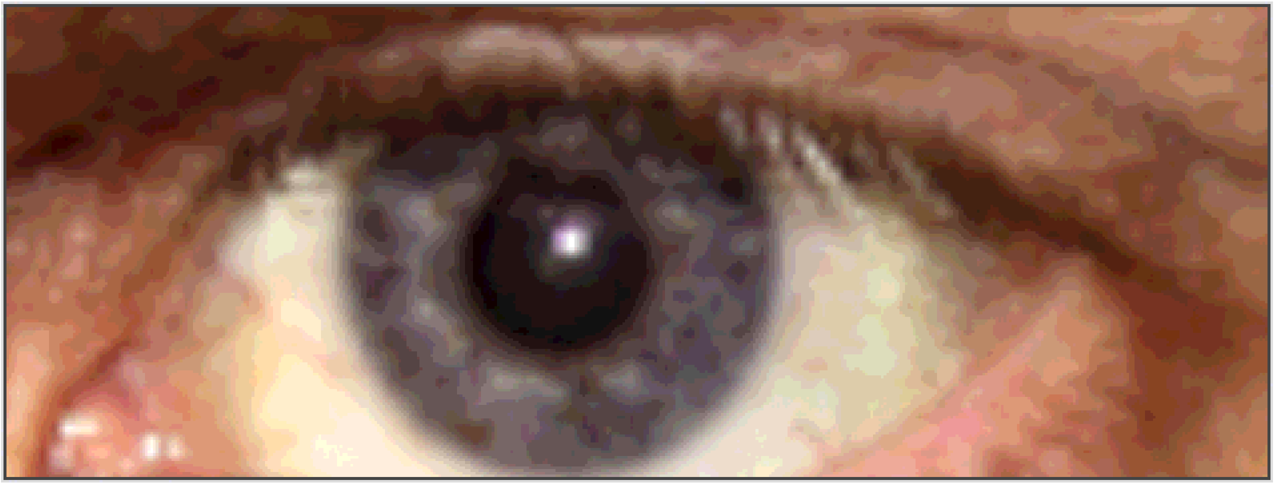

Bilinear interpolation produces relatively smooth edges with hardly any jaggies. An output pixel is a bilinear function of the 4 pixels in the nearest 2-by-2 neighborhood of the pixel in the source image.

Bicubic interpolation

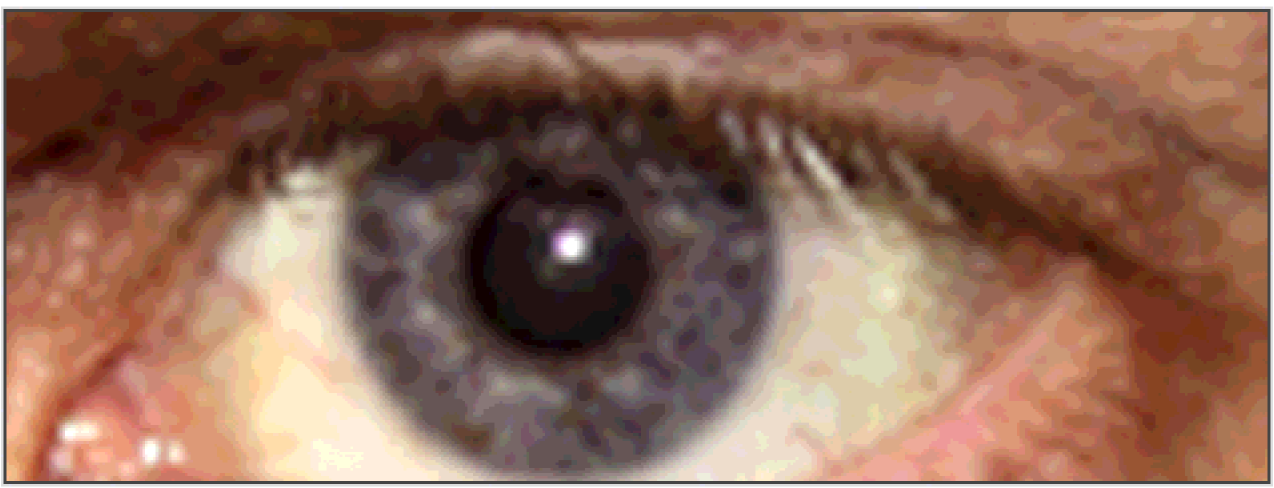

Bicubic interpolation is an even more sophisticated, it produces very smooth edges with hardly any jaggies (better than Bilinear). An output pixel is a Bicubic function of the 16 pixels in the nearest 4-by-4 neighborhood of the pixel in the source image. This is the method most commonly used by photo editing software, printer drivers and many digital cameras for up or downsizing images.

Fractal interpolation

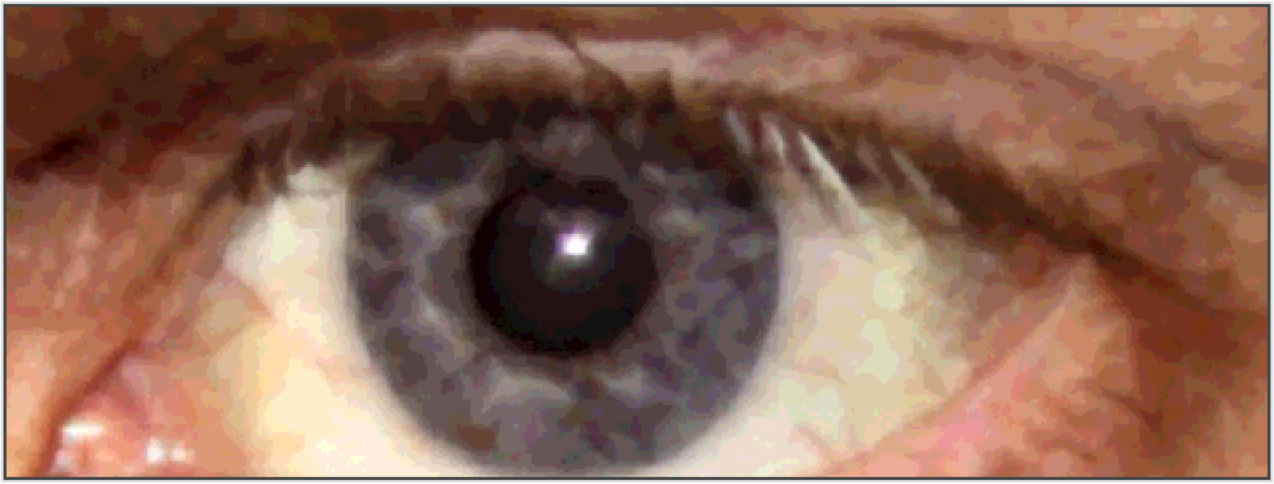

The Altamira Group have a specialized product called Genuine Fractals Pro which was specifically developed for lossless manipulation and resizing of digital images. Their algorithm produces cleaner, sharper lines when resizing digital images and is particularly good for resizing images for large prints.

There are of course many other methods of interpolation but they're seldom seen outside of more sophisticated image manipulation packages.

Interpolation as final output

Some manufacturers (notably Fujifilm with their recent SuperCCD implementation) feature their own type if interpolation to generate the native resolution image. When the Fujifilm 4700Z was first announced it was released as a 4.3 megapixel digital camera, only later did we find out that the camera had only a 2.4 megapixel CCD (and Fujifilm were forced to remove 4.3M labels from the front of the cameras) and that internal processing algorithms were processing this (interpolating) it up to the 4.3 megapixel image. While there's no doubting that Fujifilm's algorithm is far more advanced than simple Bicubic interpolation there's one rule of thumb you can't avoid: You can't create detail you didn't capture. Some digital cameras with interpolated output image resolutions:

* Fujifilm 4700Z (2.4MP -> 4.3MP)
* Fujifilm 4900Z (2.4MP -> 4.3MP)
* Fujifilm 40i (2.4MP -> 4.3MP)
* Fujifilm S1 Pro (3.2MP -> 6.1MP)
* Sony DSC-F505V (2.6MP -> 3.2MP) *
* Epson PhotoPC 3000Z (3.1 -> 4.8MP) *
* Kodak DC290 (2.1 -> 3.3MP)
* Ricoh RDC-7 (2.1 -> 7 MP) *

* These cameras offer the interpolated resolution over and above the standard high resolution, normally to make "straight to print" easier for beginners, however there's no avoiding the fact that the interpolated resolutions are mentioned in marketing material, and thus may be misleading for new users.

Sensitivity (ISO)

In traditional film photography the ISO (ASA) value of a film represents the film's sensitivity. A film with a lower ISO (e.g. 25) requires more light to create the same image than a film with a higher ISO (e.g. 800), therefore higher ISO films are suited to taking high speed (or low light) photographs. Faster films tend to be more grainy with poorer color response than slower films. Most people typically use ISO 100 or 200 (sometimes referred to as cloudy) film.

In a digital camera the sensitivity depends on the sensor (CCD/CMOS device) which, compared to film sensors are relatively "slow" devices with an optimum sensitivity of about ISO 100. In a traditional camera you put in a film of a particular sensitivity and you're then stuck with it until you finish the roll (you can change the roll in the case of APS). On modern digital cameras you can select the ISO sensitivity on the fly. This ability to quickly switch sensitivity is another attraction of digital cameras.

Unfortunately, there is a price to pay for this flexibility. A CCD is an analogue device, this means it outputs a certain voltage (normally VERY small) for a certain amount of light which is subsequently digitized by the analogue to digital converter, when you increase the sensitivity you're really just turning up the amplification of this signal (like the volume control on a stereo amplifier), the trade off is that you also amplify the "dark current" (noise) and so higher ISO images from digital cameras often exhibit noise.

Some recently advancements in sensor technology have improved the situation with many higher end digital cameras capable of sensitivity above ISO 400 without impacting on image quality too greatly.

Sharpening

Sharpening in a digital sense is the enhancing of edge detail. It is performed by a mathematical formula that is applied across the image. Put simply it enhances the visibility of a boundary between light and dark tones in an image.

In-camera sharpening

Most consumer digital cameras will, as a part of normal image processing apply some level of sharpening, to counteract the effects of the interpolation of colors during the color filter array decoding process (which will soften detail slightly). The problem with in-camera sharpening is that it increases the visibility of jaggies and can increase the visibility of other image artifacts.

Recently, digital cameras have allowed users to control the amount of sharpening applied to an image (menu option), the primary reason for this is the widespread use of image processing applications, the "digital darkroom" and the ability to control artifacts related to sharpening. Higher-end digital SLRs and other cameras which support RAW format don't apply any sharpening to the RAW image, this allows the user to decide on the level of sharpening when acquiring the image later.

Canon took a stance on sharpening with their new EOS-D30, the images straight out of the camera have almost no sharpening applied, this produces very clean, artifact free images which some may consider "soft" but the consensus of opinion has been that this is probably (at least for a digital SLR) more preferable as it gives you the option to sharpen the image later if required.

Image editor sharpening

Quite a few digital camera owners choose to sharpen images even further, often to "pull out" fine details of reduced size images for web or monitor viewing purposes (sharpening for printing isn't really recommended and can make the image look fake and over-processed). Probably the most popular sharpening method is the "Unsharp Mask" (a very confusing name) which produces the most pleasing sharpening results without making the image look over-processed. It also allows for a wide variation of parameters to be used, which increases the ability to match the sharpening level to the particular image / desired output.

The samples below were sharpened in Photoshop 6.0 with an Unsharp mask at different percentages, radius 0.6 pixels, threshold 2 levels. (I personally don't like to use an Unsharp mask at a radius above 0.8 pixels as this will tend to introduce the dreaded "halo").

No sharpening

Unsharp Mask 83%


Unsharp Mask 120%


Unsharp Mask 500%

White Balance

White balance is a name given to a system of color correction to deal with differing lighting conditions. Normally our eyes compensate for different lighting conditions, but when taking a still with a digital camera the camera has to find the "white point" (the assumption that a white object must appear white) to correct other colors cast by the same light.

Most digital cameras feature automatic white balance, this means that the camera looks at the overall color of the image and calculates the best-fit white balance, however these systems are often fooled (especially if taking a photograph dominated by one color, say green). Most digicams also allow you to override the automatic white balance by choosing a white balance manually, typically sunlight, cloudy, fluorescent, incandescent etc.

Modern "prosumer" digital cameras also allow "white preset" which simply means measuring the white point from a white sheet of paper or card (or nearby wall), the camera will then record that temperature and use it to correct all images until you reset it.

Color Temperature

Each type of light can also be represented by a numerical color temperature; here are the (rough) color temperatures of typical lighting conditions:

Type of light	**Color temperature**
Incandescent	2500K - 3500K
Twilight	4000K
Fluorescent	4000K - 4800K
Sunlight	4800K - 5400K
Cloudy daylight	5400K - 6200K
Shade	6200K - 7800K

Photographic

Aperture

In a camera system the aperture is an adjustable diaphragm of over overlapping blades which can be thought of like the iris of the eye. The aperture value represents a ratio of the equivalent focal length of a lens to the diameter of its entrance pupil. An aperture can be represented in several different ways: f/8, F8, 1:8 (each are equal).

f/# = f /A

Where: f = lens focal length, A = diameter of aperture opening.

The larger the f-number the smaller the aperture, and thus the less light it allows in. Each step of in the 'normal' aperture range allows exactly half the amount of light (e.g. F2.0, F2.8, F4.0, F5.6, F8.0, ...). Aperture also has a direct effect on depth of field and in conjunction with a metered value in automatic exposure see also shutter speeds.

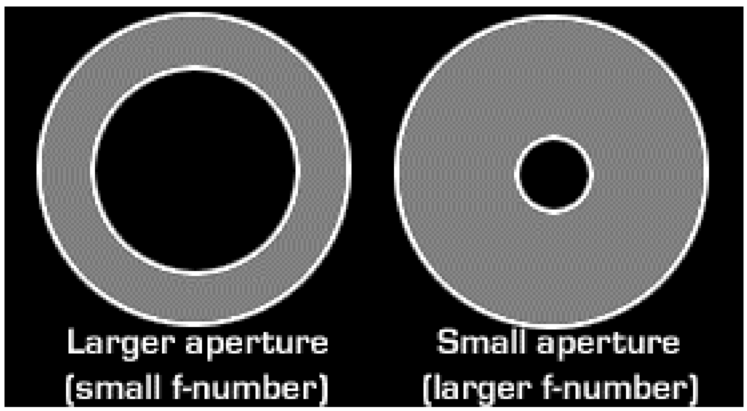

You will also note that lenses are also marked with their "maximum aperture" (zoom lenses with maximum at both wide and tele focal lengths), this represents the effective aperture that the glass within the lens causes to the light passing through it, your largest possible aperture when using this lens.

Lenses with a larger maximum aperture (e.g. f/2) would be called fast (they allow more light through to the sensor).

Auto Bracketing

Bracketing is a technique used to take a series of images of the same scene at a variety of different exposures that "bracket" the metered exposure (or manual exposure). "Auto" simply means the camera will automatically take these exposures as a burst of 2, 3 or 5 frames with exposure settings of anything between 0.3 and 2.0 EV difference. This can be useful if you're not sure exactly how the shot will turn out or are worried that the scene has a difficult dynamic range. On a digital camera this can also be used to combine under and over exposed images together to produce an image with more dynamic range than the camera can normally capture.

When setting up for bracketing you can usually select the number of frames to be taken (typically 3 or 5), the exposure setting and the order in which to take the shots (e.g. 0, -, + or -, 0, + etc.). The example below (very subtle) was taken with auto bracketing of 3 frames at 0.7EV in the -, 0, + order. Thus, in this case without bracketing the camera would simply have shot the frame as 1.3s @ F14.


1.0s, F14 (-0.3 EV)


1.3s, F14 (0.0 EV)


1.6s, F14 (+0.3 EV)

Exposure/EV

The exposure is the actual exposure of the digital camera's sensor (CCD/CMOS) to light when taking a photograph. Most modern digital cameras used a mechanical shutter (just like traditional cameras) to control the exposure of the sensor, some use electronic shutters, an electronically timed measurement of sensor output. In film photography the name refers to a negative's exposure to light.

A value is given to a single exposure; it's Exposure Value or EV. This is the combination of the sensitivity of the CCD/CMOS (or ISO), aperture and shutter speed. Thus a single "exposure" could be ISO 100, f/2.4, 1/60s. Fully automatic digicams (all at the moment) will calculate all of these values for you, some of the more modern cameras have manual features which allow you to set the aperture (aperture priority) or the shutter speed (shutter priority) and even the ISO sensitivity of the CCD.

In a true photographic sense, Exposure Value (EV) is a single number which refers to the amount of light for a given exposure, this value is then used to calculate the correct combination of aperture and shutter speed (at a given sensitivity), the table below represents the relationship between shutter speed, aperture and EV at ISO 100.

	Shutter Speed													
Aperture		1	1.4	2	2.8	4	5.6	8	11	16	22	32	45	64
	1 s	0	1	2	3	4	5	6	7	8	9	10	11	12
	1/2 s	1	2	3	4	5	6	7	8	9	10	11	12	13
	1/4 s	2	3	4	5	6	7	8	9	10	11	12	13	14
	1/8 s	3	4	5	6	7	8	9	10	11	12	13	14	15
	1/15 s	4	5	6	7	8	9	10	11	12	13	14	15	16
	1/30 s	5	6	7	8	9	10	11	12	13	14	15	16	17
	1/60 s	6	7	8	9	10	11	12	13	14	15	16	17	18
	1/125 s	7	8	9	10	11	12	13	14	15	16	17	18	19
	1/250 s	8	9	10	11	12	13	14	15	16	17	18	19	20
	1/500 s	9	10	11	12	13	14	15	16	17	18	19	20	21
	1/1000 s	10	11	12	13	14	15	16	17	18	19	20	21	22
	1/2000 s	11	12	13	14	15	16	17	18	19	20	21	22	23
	1/4000 s	12	13	14	15	16	17	18	19	20	21	22	23	24

Another way of calculating the exposure value is this formula:

EV = log2(aperture2 x (1/shutter speed) x (ISO sensitivity/100))

Exposure Compensation

Exposure Compensation is the ability to override the cameras metered exposure by a preset value normally in the range of -2 to +2 EV in 0.3 EV steps. For instance if your camera metered a scene at 1/60s F5.6 (11 EV) and you had an exposure compensation of +0.7 EV it would actually take the shot with an exposure of 1/40s F5.6 (10.3 EV), yes inverse of what you'd expected but it would be even more confusing for +VE to produce darker images and -VE to produce lighter images.

This is useful once you understand your cameras metering system and where and when you should override it. For example most Pro's recommend using at least -0.3 or -0.7 EV compensation for over-the-shoulder bright sunlight shots and +0.7 or +1.0 EV compensation for strongly back-lit shots.

Flash Output Compensation

Flash output compensation is similar to exposure compensation in that it allows you to preset an adjustment value for the flash output power. Some digital cameras allow you to set this value using the familiar EV range (+/-2EV); others simply have a "high, normal and low" setting. Whichever way it's often useful as it allows you to compensate for situations where either (a) the cameras flash metering isn't perfect or (b) a difficult scene that requires some flash compensation. See also exposure, metering.

Manual

Manual controls on a digital camera are becoming more common and desirable as traditional photographers previously uninterested by low resolution digital cameras begin to move into the digital realm and demand more control over the camera.

See the following articles:

* Manual Focus
* Aperture Priority
* Shutter Priority
* AE Lock
* Full Manual

Manual Focus

Manual focus (on consumer grade digital camera's) is the ability to disable the cameras built in automatic focus system and focus the lens by hand, in digital cameras it is normal for manual focus to be implemented on a fly-by-wire basis, that is that inputs to focus in or out are relayed to the autofocus system which effects a change in focus. Manual focus becomes important in low light, macro or special effect photography; it can also be very useful when the autofocus system is unable to get a good focus lock. Some digital cameras allow you to manually focus only to a few preset distances.

More sophisticated digital cameras (notably the professional end) feature the same focus systems as their inherited camera bodies and therefore allow focusing using the normal focus ring on the attached lens.

Aperture Priority

Aperture priority is the ability to set the aperture over its full range and have the camera calculate the best shutter speed to expose the image correctly. This is important if the photographer wants any control of depth of field or for special effects.

Shutter Priority

Shutter priority is the ability to set the shutter speed over its full range and have the camera calculate the best aperture to expose the image correctly. Shutter speed priority is often used to create special effects such as blurred water on a river / waterfall or to freeze action in sports scenes.

Full Manual

Full Manual exposure is the ability to set both the aperture and shutter speed. This gives the photographer ultimate control over the exposure and is very useful for ensuring the same exposure is used for a sequence of shots. Higher-end prosumer digital cameras and all digital SLRs feature full manual exposure, when in full manual exposure mode the camera will often display a simulated exposure meter which will indicate how far over or under exposed the image is compared to the exposure value calculated by the cameras metering system.

Prosumer digital cameras with live LCD preview will often simulate the effects of the exposure on the live preview.

AE Lock

Automatic Exposure lock is the ability to lock the cameras calculated exposure settings (aperture and shutter speed) over a series of images, very important for producing panoramas, when stitching images together each image should have the same exposure.

Depth of field

Depth of field is a term which refers to the areas of the photograph both in front (closer to you, about 1/3 of the DOF) and behind (further away, about 2/3 of the DOF) the main focus point which remain "sharp" (in focus). Depth of field is affected by the aperture, subject distance (closer subjects produce a shallower depth of field), and focal length (shorter focal lengths produce larger DOF, thus a 28mm lens at f/5.6 produces a greater depth of field than a 70mm lens at the same aperture).

Put simply, a larger aperture (smaller f-number, e.g. f/2) has a shallow depth of field, anything behind or in front of the main focus point will appear blurred. A smaller aperture (larger f-number, e.g. f/11) has a greater depth of field; objects within a certain range behind or in front of the main focus point will also appear sharp.

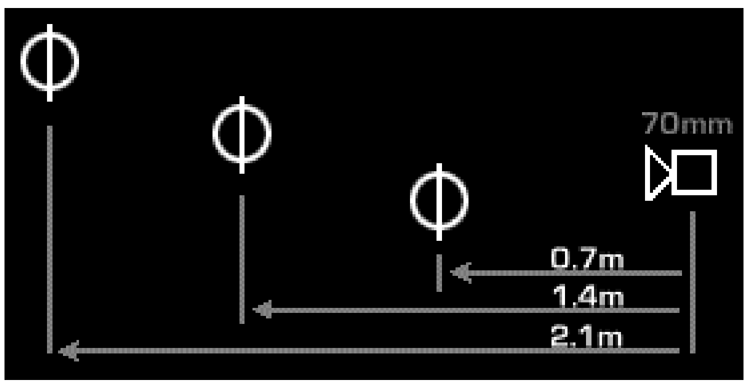

This is the setup that was used to produce the example below, three postcards 0.7m apart, the camera; a Canon Pro 70 set at its maximum telephoto (70mm) focused on the first card.

As you can see at a large aperture of f/2.4 only the first card is in focus, at f/8 the middle card is sharp and the distant card is almost sharp.

There is an added technique you can use to gain more depth of field, you'll note that above I quoted that 1/3 of the depth of field is in front of the focal point and 2/3 is behind, knowing this you could focus "in front" of the main subject (closer to you) and still render them as sharp as they fall into that large 2/3 of the DOF. The full SIZE of the DOF depends on the attributes noted above (aperture and subject distance).

Shutter Speed

The shutter speed is the length of time the "shutter" which covers the CCD allows light onto the CCD, in older digital cameras this was an electronic shutter (electronically start and stop recording light), now being replaced by more traditional mechanical shutters (explanation below).

A shutter speed of 1/125s means that the sensor is exposed to light for 1/125s, the resulting light levels (which are output as voltages by the sensor) are recorded for each pixel of the image. Most people can hold a camera still enough not to get shake at shutter speeds of 1/60s and above, below that speed a tripod, monopod, physical object or optical stabilization. Most prosumer digital cameras have a fairly wide range of shutter speeds from relatively long (8/16 seconds) up to about 1/1500 sec. Some prosumer cameras provide shutter priority exposure modes.

Typically if you're taking photographs where you need to "freeze" the action you'll need a shutter speed of 1/250s. To avoid shake you should always have shutter speed faster than the focal length of the lens (or current zoom) in mm; for example a 300mm lens would need a shutter speed of 1/300s to avoid blurring (due to camera shake).

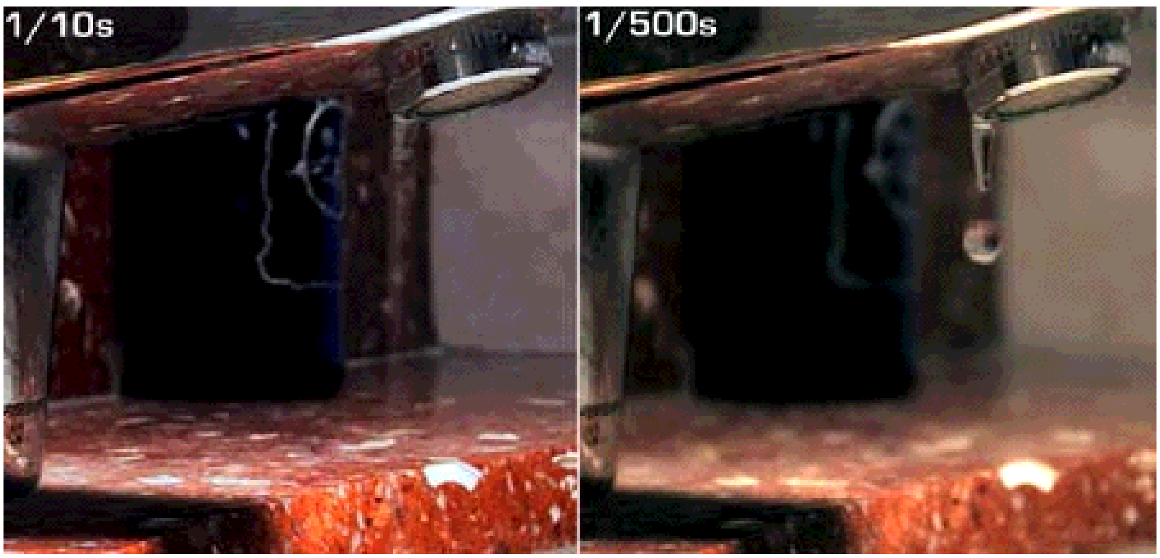

In this example the image on the left was taken with a shutter speed of 1/10s, the drips of water coming from the tap are blurred at this speed (you'll need to click on the image to view the larger version to appreciate this), on the image on the right was taken at 1/500s, the droplet of water is frozen in mid-air (note this image had to be brightened in Photoshop because of the poor lighting available for the shoot).

Mechanical vs. Electronic shutters: The electronic shutter is simply the cameras onboard electronics taking a "snapshot" of light for the shutter speed period of time. A mechanical shutter physically blocks the light to the CCD before exposure opens for the shutter speed period of time and then covers the CCD over again (just like a film camera). More sophisticated digital cameras will even take a "snapshot" (very quickly) just before the exposure to get a black-level which can be used to subtract from the taken image to reduce noise.

Metering

The metering system in a digital camera is the system which measures the amount of light (EV value, exposure value) in the current frame and calculates the best-fit (depending on metering mode - below) exposure. Automatic exposure is standard in all digital cameras, and all you have to do is select the metering mode, point the camera and press the shutter release, nine times out of ten the camera will expose the image pretty much correctly.

The metering method defines how and what part of the image is used to calculate the exposure value, different manufacturers offer different metering modes but there are still similarities between these implementations and we've summarized a few of these below:

Center-Weighted Average Metering

Probably the most common metering method implemented in nearly every digital camera and the default for those digital cameras that don't offer metering mode selection. This method averages the exposure of the entire frame but gives extra weight to the center.

Spot (Partial) Metering

Spot Metering allows you to meter for the subject directly in the center of the frame (or on some cameras at the selected AF point), this is normally a small circle just 10% the area of the whole frame, the exposure of the rest of the frame is ignored. Useful for brightly Bakelite or macro shots.

Matrix or Evaluative Metering

This is probably the most complex metering mode, and each manufacturers exact detail on how they implement matrix metering is closely guarded. Sufficient to say that the scene is split up into a matrix of metering zones (between 30 and 200), each of which are evaluated individually, the overall exposure is based on an algorithm specific to that camera. Matrix metering often offers the best exposure possible.

Remote Capture

Remote Capture normally refers to software which is capable of remotely firing the camera (normally tethered) to take images under the control of a computer. Recent examples of remote capture software can be seen supplied with Fujifilm's S1 Pro and Canon's EOS-D30 which are both remotely controlled over the USB port. Remote capture is useful for two reasons (a) images can be stored directly onto the computers hard disk and (b) in a studio environment images can be immediately previewed on a larger screen / LCD display.

Time Lapse

Time Lapse simply means a number of frames shot automatically by the camera over a period of time or with a certain time interval between each frame. Example: This mode could be used to set up the camera on a tripod and take (say) a frame an hour of a flower opening. Some cameras feature a built-in Time Lapse mode; others allow you to set up Time Lapse as part of a Remote Capture application, and this requires the camera to be tethered to a computer.

Optical

Barrel Distortion

Barrel distortion is a lens effect that causes images to be spherized at their center. Barrel distortion is associated with wide-angle lenses and only occurs at the wide end of a zoom lens. Most noticeable when you have a very straight edge near the side of the image frame. Some people find this to be an unacceptable fault of the camera; my own personal feelings are that although barrel distortion can be visible sometimes it would not be noticeable in 90% of your photography.

We measure barrel distortion in our reviews as the amount a reference line is bent as a percentage of picture height, for most consumer digital cameras this figure is normally around 1%. See also the opposite effect, pincushion distortion.

Chromatic Aberrations

Also known as "purple fringing". Well, maybe and maybe not. It's certainly something we recognize as being a problem on consumer digital cameras, how many times you'll see it in a real life shot really depends on how you shoot and if your careful or not. I know of many Olympus C-3030Z owners who are very happy with their cameras and never let the fact that this camera exhibits some of the worst chromatic aberrations we've seen in our reviews. (Examples of chromatic aberrations below).

Olympus C-3030Z

Canon PowerShot G1

The description of Chromatic Aberrations below is purely from an optical point of view, what often happens in digital cameras is that the visibility of chromatic aberrations are amplified by an effect called blooming which involves the overflow of charge from oversaturated pixels to neighboring pixels on the digital camera's sensor.

Chromatic Aberrations in a single lens

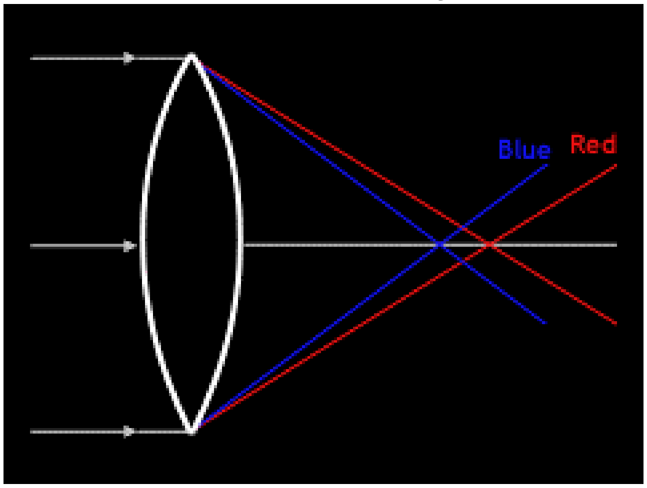

Put simply, chromatic aberrations are by the cameras lens not focusing different wavelengths of light onto the exact same focal plane (the focal length for different wavelengths is different). The amount of chromatic aberration depends on the dispersion of the glass.

Achromatic / Apochromatic Doublets

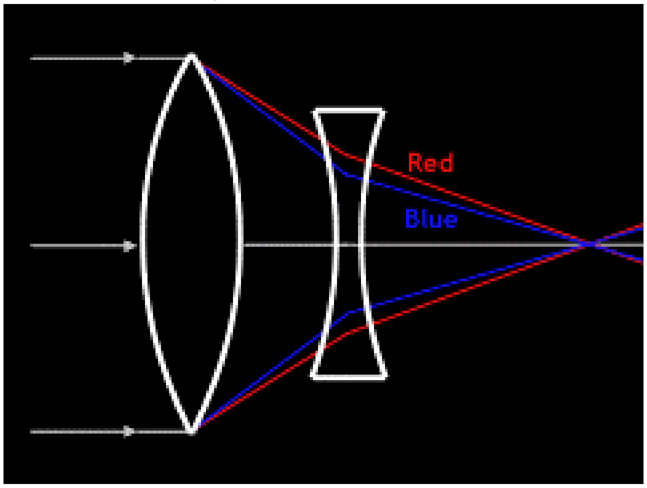

Special lens systems (achromatic or apochromatic doublets) using two or more pieces of glass with different refractive indexes which reduce or eliminate this problem, but not even these lens systems are completely perfect and (especially it would seem at wide angles) still can lead to visible chromatic aberrations.

Repairing Images with Chromatic Aberrations

It is possible to at least decrease the visibility of chromatic aberrations in images, it requires altering the saturation of the specific magenta hue generated by the chromatic aberrations, this doesn't completely fix the problem because it simply replaces the magenta with a gray shade but it certainly makes it much less visible. The technique below is described using Photoshop 6.0 (though it should also work in 5.0/5.5).

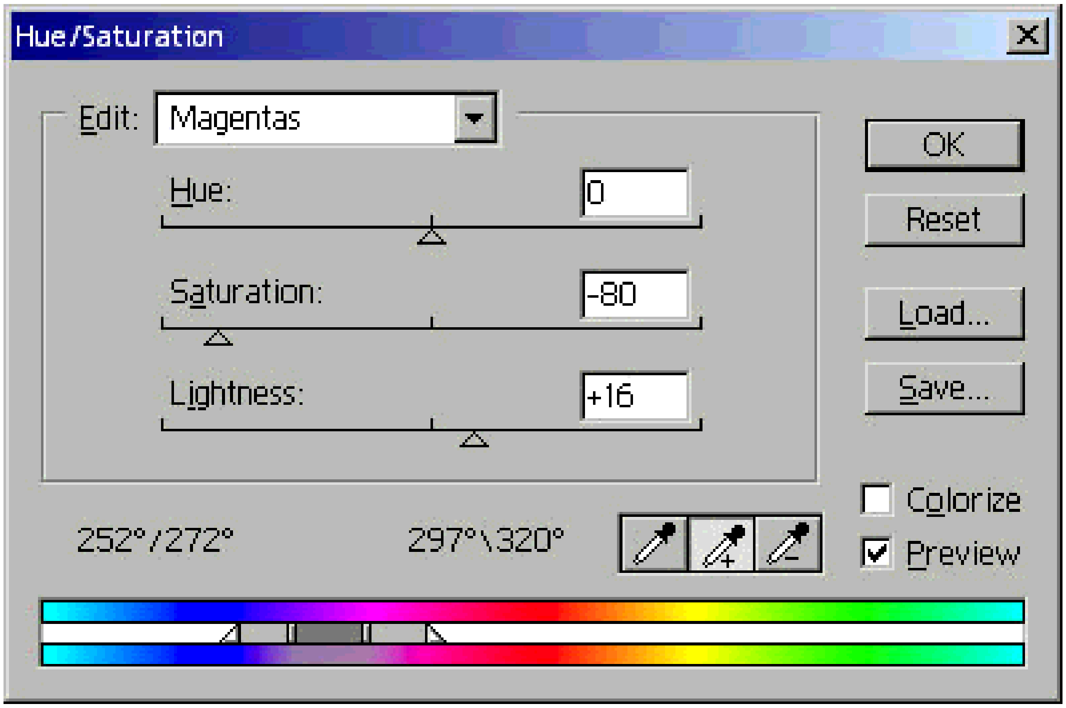

1. Open the image to be repaired
2. Zoom to 100% so that an area of the image with chromatic aberrations is visible
3. Select Image -> Adjust -> Hue/Saturation (PC: CTRL+U, Mac: Apple+U)
4. Change the Edit: dropdown to "Magentas", the spectrum display should now show the color selector / ramp bars.
5. Change Saturation to -80 and Lightness to +16 (you may wish to play with these figures)
6. Now gently push the color ramps up and down the spectrum around the magenta level until you turn all the chromatic aberration completely gray. Click OK.
7. That's it.


Before repair

After repair

Focal Length

Focal length is measured in millimeters, it is defined as the distance from the lens to a point where parallel rays are focused to a point (diverge), traditionally measured in millimeters (mm). The lens on a consumer digital camera is marked with its focal length (or range of focal lengths if it's a zoom lens). This is typically a very small number such as 6 - 15mm, in traditional 35mm photography everyone is "used to" the common focal lengths of 28mm, 50mm, 200mm etc.

Because the sensor in a digital camera is much smaller than a 35mm negative the lenses can be made smaller (because of this they have to be of a much higher quality), to get the true focal length you need to multiply this small size by a value called the "focal length multiplier" (this is especially important for digital SLRs which take normal 35mm lenses).

For example, the lens on the Canon PowerShot Pro 70 is marked 6 - 15mm, however the "zoom" of the lens is always quoted as 28mm - 72mm, therefore to make the focal lengths more understandable we are actually multiplying the lenses focal length by 42/3.

35mm Equivalent

All digital camera manufacturers publish this "35mm equivalent" focal length simply because people are used to hearing it and knowing what kind of image a 28mm lens produces compared to a 50mm lens, some common lens sizes (focal lengths):

* <20mm = Super Wide angle
* 24mm - 35mm = Wide angle
* 50 mm = Normal, the same picture angle as your eye (46o).
* 80mm - 300mm = Telephoto
* >300mm = Super Telephoto

Please note also that focal length has a direct effect on perspective.

Optical zoom (x times zoom)

You may also see the abilities of the cameras lens quoted as a "zoom", e.g. 2 x zoom. This is really a bit of a throw back to video cameras; all the zoom describes is the lenses ability to multiply the size of a subject between its minimum and maximum focal lengths. For example a person standing directly in front of you may fill half the frame height when the lens is at 35mm, if you 'zoomed in' to 70mm (2 x 35mm) they would take up the whole frame height.

Some digital cameras suffer from barrel distortion at small focal lengths (wider angles: <35mm).

Focal Length Multiplier

A Focal Length Multiplier applies to 35mm lenses used on digital SLRs with sensors smaller than the normal 35mm negative frame. The focal length multiplier is expressed as a scale factor such as 1.5x or 1.6x, thus on a digital SLR with a focal length multiplier of 1.5x a 28mm lens would produce the equivalent picture angle (strictly "field of view") of a 42mm lens. Think of the sensor in a digital SLR as simply "cropping" a middle portion out of the 35mm sized frame projected by the lens.

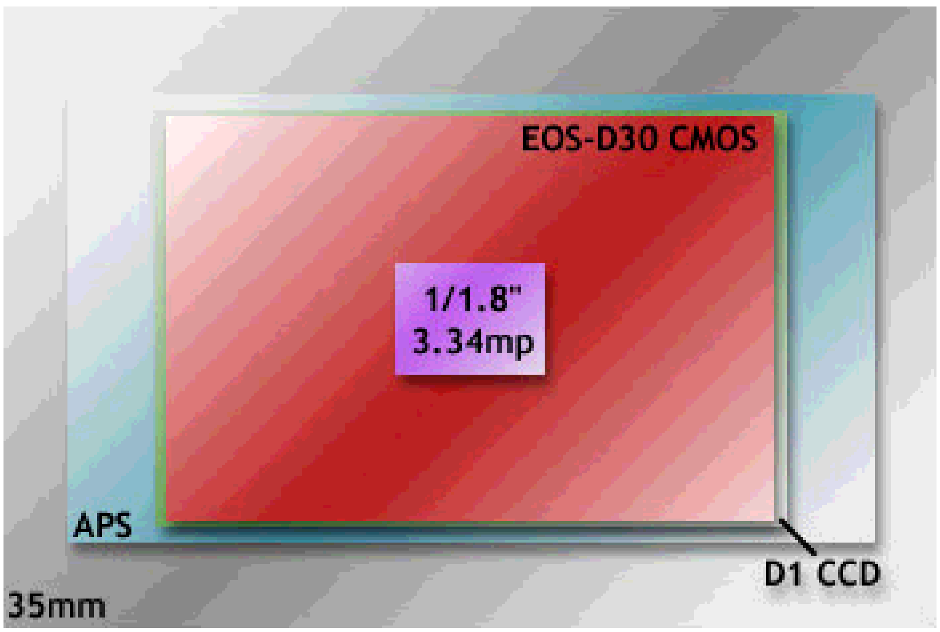

Interestingly, despite the narrower picture angle and effective focal length of 42mm, the perspective of objects in a scene taken with this lens remains the same (only changing the subject distance would affect the perspective) as it would as a 28mm on a 35mm film SLR. Examples of focal length multipliers below.

Camera / Sensor	**Sensor size (mm)**	**Horiz. size as % of 35mm**	**Focal Length multiplier**
Sony 1/1.8" CCD *	5.52 x 4.14	15%	N/A
Nikon D1 / CCD	23.6 x 15.5	67%	1.48x
Canon EOS-D30 / CMOS	22.0 x 14.9	65%	1.59x
35mm negative	35.0 x 23.3	100%	1.00x

Image Stabilization

For some time manufacturers of zoom or telephoto lenses for SLR cameras and high-end binoculars have offered an implementation of image stabilization, it has also been available in digital video cameras with large zooms. Image stabilization helps to steady the image projected back into the camera by the use of a "floating" (often connected to a fast spinning gyroscope) optical element which helps to compensate for high frequency vibration (hand shake for example) at these long focal lengths. Canon EF SLR lenses with image stabilization have an IS suffix after their name, Nikon uses the VR "Vibration Reduction" suffix on their image stabilized Nikkor lenses.

It's said that image stabilization can help you take handheld shots almost two stops slower than with image stabilization off. More recently digital cameras with large zoom lenses have also featured image stabilization. Examples of digital cameras with image stabilization: Sony FD-91, 95, CD1000, Olympus C-2100UZ, E-100RS.

Important footnote: **Only Optical Stabilization** will work for digital still cameras, digital image stabilization (seen on some "lesser" digital video cameras) works by pixel shifting the image (not something that will work for a digital still camera).

Macro

In strict photographic terms the word Macro means the optical ability to produce a 1:1 or higher magnification of an object on the film negative, that is get very close (not always physically) to a very small object (a bit like a microscope does, but obviously not as much). The second type of lens is a close-up lens, normally anything less than 1:2 is seen as close-up (rather than Macro).

On digital cameras there is often a Macro Focus mode, this switches the auto focus system to attempt to focus on subjects much closer to the lens to take interesting shots of small objects. The macro abilities of a digital camera vary, and most should be ashamed of calling the focus mode "macro" (it should strictly be called "close up"). We measure macro ability (of cameras with non-interchangeable lenses) in our reviews as the ability of the lens to get the best possible frame coverage.

The clear outright winner (without add-on lenses) are Nikon's Coolpix 950 and 990 both of which can produce amazingly close shots which can clearly be called macro. The 990 for example can get close enough to fill it's 2048 x 1536 frame with just 0.7" of an object.

Lenses

Most consumer digital cameras have non-interchangeable zoom lenses that have been designed to work with a specific sensor size. Quite a few "prosumer" camera manufacturers also have add-on lens adapters to either expand the picture angle (wider) or narrow it (telephoto). These lenses are often labeled with a focal length multiplier such as x0.8 for a wide-angle adapter (eg. 35mm becomes 28mm) or x2.0 for a telephoto adapter (e.g. 115mm becomes 230mm). These adapters often can't be used across the whole range of a zoom lens because they would introduce vignetting (the barrel or sides of the lens become visible) at the opposite end.

Lens Quality

Because of the small size of a digital camera's sensor the lenses used in digital cameras have to be of much better quality than glass which would be "acceptable" on a 35mm camera. Digital SLRs with larger sensors (at least at the current 3 megapixel size) haven't yet run into problems with lower resolution 35mm lenses.

Perspective

Perspective is an effect caused by the picture angle that the focal length of the lens produces, it changes the way an image looks both in the size of objects and depth of view in the image. At wider angles (read smaller focal lengths, e.g. 28mm) the background appears to be MUCH further away from any subject in the foreground than your eye would normally see (exaggerated perspective) and narrow angles (reader longer focal lengths, e.g. 200mm) the distance between background and foreground objects seems much shorter (compressed perspective).

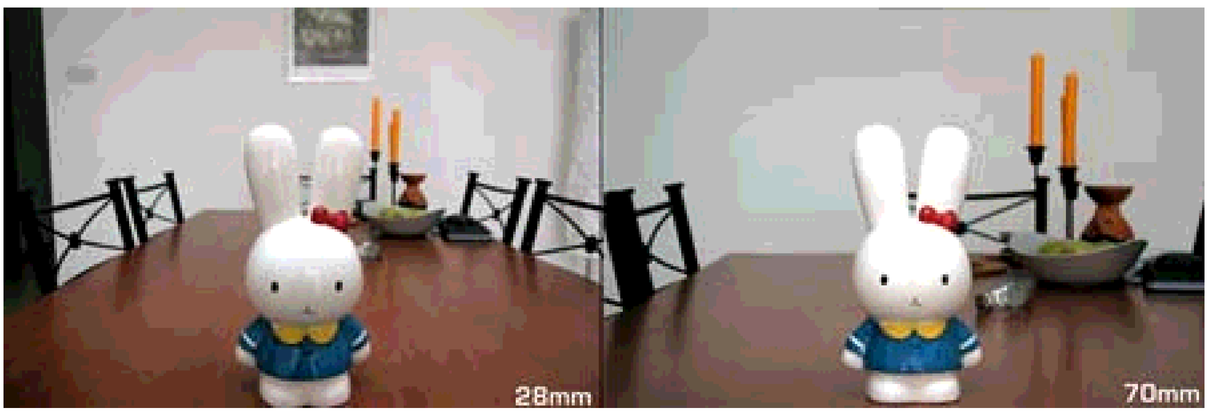

In the above example the subject was not moved, and is (approximately) the same height in each image, however in the first image taken with the lens set at 28mm the table appears to be very long, candles and bowl of fruit in the background appear to be very far away (exaggerated perspective). On the right we see the same scene taken at 70mm and everything now seems much closer together (compressed perspective) and objects in the background and foreground are more similarly sized.

In the above example the position of the camera had to be changed subject distance to keep the size of the primary subject the same.

Picture Angle

Put simply the picture angle is the field of view the lens produces, the angle in degrees of view from the lens out to the scene. A smaller focal length (such as 28mm) produces a wide picture angle (wide angle), a larger angle of view. A larger focal length (such as 200mm) produces a narrower picture angle (narrow angle). A 50mm lens is called a normal lens because it produces roughly the same picture angle as the human eye (about 46 degrees).

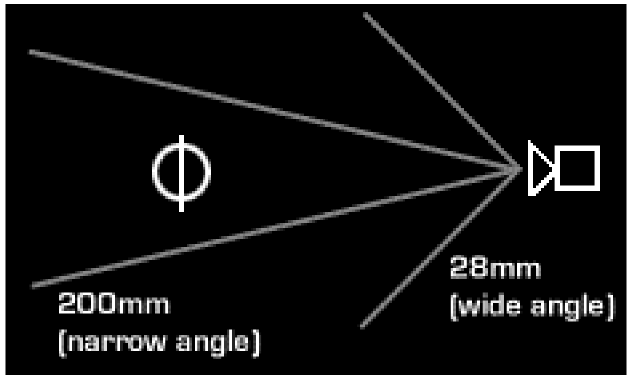

Pincushion Distortion

Pincushion distortion is a lens effect that causes images to be pinched at their center. Pincushion distortion is associated with zoom lenses or when adding telephoto adapters and only occurs at the telephoto end of a zoom lens. Most noticeable when you have a very straight edge near the side of the image frame.

We measure barrel distortion in our reviews as the amount a reference line is bent as a percentage of picture height, for most consumer digital cameras this figure is normally less than for barrel distortion at around 0.6%. See also the opposite effect, barrel distortion.

Subject Distance

Subject distance is just the distance from the camera (lens) to the main subject, varying this distance along with the focal length will produce different perspectives.

Varying the subject distance with the same aperture will produce different a depth of field.