A Technical Summary of Film Photography

Anthony G. Atkielski

Photography is the science and art of recording visual images of the real world in a permanent and tangible form, using a device called a *camera* to partially automate the process. The purpose of this paper is to provide a concise overview of the *science* of photography (as opposed to the art, which is a separate subject), as seen from the photographer's standpoint. We will discuss only film photography, as it is still largely predominant over digital photography, particularly among serious amateur and professional photographers. Furthermore, although there are many types of film photography, we will concentrate here upon the most popular type, namely, so-called *35-millimeter* photography.

The Essentials

As we have described above, photography uses a *camera* to record images. All photographic cameras can be divided into three parts: (1) a replaceable strip, sheet, or plate, usually made of plastic and called *film*, that is coated with a blend of chemicals that undergo an invisible reaction when exposed to light; (2) a box, called a *camera body*, that holds the film and protects it from light except during the instant required to take a photograph; and (3) an optical device, called a *lens*, that can project and focus reflected light from an image onto the film. Each of these parts of the photographic process will be described in greater detail later in this paper.

It is important to note the distinction between the *camera*, which is the entire device used to take pictures, and the *camera body*, which is a component part of the camera that holds the film. In most cameras used by professionals and serious amateurs, the three main parts of the camera (film, camera body, and lens) are interchangeable; that is, any one of several lenses can be attached to the camera body, and any one of several types of film can be loaded into it.

Figure 1

Figure 1 shows a typical, professional photographic camera. The lens is the cylindrical structure on the front of the camera; the camera body is the squarish structure behind the lens. The film is concealed within the camera body to protect it from light and is not visible in this illustration.

The full process of photography also includes one or two additional steps, called *development* and *printing*, to produce a usable image, but since these steps are normally handled by a commercial film laboratory, rather than by the photographer himself, they will be only peripherally touched upon in this paper.

Taking a Photograph

In order to capture an image photographically, it is necessary to project and focus light reflected from the image onto a sheet of light-sensitive photographic film, using a lens made of high-quality optical glass. When the light from the image is projected onto the film, the film undergoes an invisible chemical change caused by absorption of the energy contained in the light. The greater the exposure to light, the greater the chemical change; that is, bright portions of the image undergo more of a change than dark portions. These changes create what is called the *latent image*: a latent image is chemically present in the film, but it is not visible to the eye. The process of capturing the image on film in this way is called *exposure*, and film that contains latent images is said to be *exposed*.

Because the effects of light upon the film are cumulative and irreversible, a given portion of the film in a camera can be exposed only once. The camera body protects the film from light except during the exposure, and moves the film so as to present a fresh, unexposed section of film for exposure for each photograph taken.

Since real-world images vary enormously in the amount of light they reflect, there must be a way to control the amount of light reaching the film from the subject image during exposure. Two camera mechanisms control this, which we will briefly describe now in order to make the overall process of taking a photograph clearer (more detailed explanations of the mechanisms will be provided later on in this paper).

The Aperture

The first of these mechanisms, called the *aperture*, is an opening inside the lens through which light must pass in order to reach the film. The size of the aperture can adjusted to vary the amount of light exposing the film. When shooting photographs in very bright lighting conditions, the aperture can be made extremely small to allow only a small portion of the light entering the lens to reach the film; conversely, in very dark conditions, the aperture can be completely opened so that all of the light entering the lens (or nearly so) reaches the film.

The Shutter

The second mechanism, called the *shutter*, is a mechanical door, located within the camera body, that completely blocks the path of light from the lens. The shutter is kept closed at all times except during an actual exposure. During an exposure, the shutter opens very briefly and very quickly in order to allow light to reach the film. The amount of time during which the shutter remains open can be adjusted to allow more or less light to reach the film. For bright subjects, the shutter is opened for only an extremely brief period; for dark subjects, the shutter is opened for a much longer period. In most photographic situations, the shutter opens for only a small fraction of a second, since most photographic film is very sensitive to light.

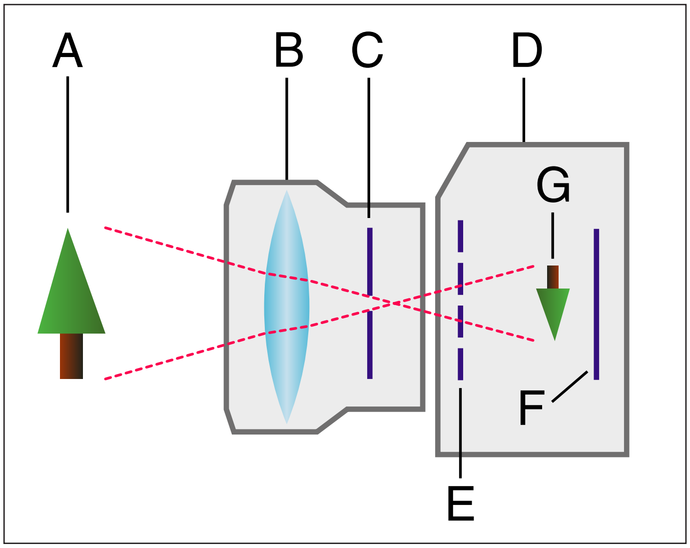

Figure 2

Taking the Photograph

The interaction of all these various components of image capture in photography is illustrated very schematically in Figure 2. Light reflected from the real-world object *A* (a tree, in this case) is gathered and focused by the optical lens *B* onto the film *F* contained in the camera body *D*. The dotted lines show the path of the light rays through the lens and into the camera. Note that the light passing through the lens is constrained by the aperture *C*, and that the passage of the light to the film within the camera body can be completely blocked by the shutter *E*.

(Although the figure shows a dotted line to represent the shutter, in fact the shutter is either completely open or closed most of the time.)

The image of the tree is upside down as it is projected onto the film because a characteristic of optical lenses is that they invert the images they project; this is of little consequence because the film can simply be turned around after development to restore the image to its correct orientation.

This is the general process of taking a photograph. We shall now look at each of the components of the camera in greater detail.

The Lens

The optical lens is the most important part of the camera, because the quality of the lens has more of an influence on the quality of the final photograph than any other aspect of the image-capture process.

Most cameras have interchangeable lenses; that is, the lens is separate from the camera body, and the former can be easily mounted or dismounted from the latter. Some inexpensive cameras incorporate a single lens permanently into the camera body. In this paper, we are concerned mostly with the 35mm cameras used by professionals and serious amateurs, all of which have interchangeable lenses, and so we will assume that the lens and camera body are separate unless otherwise indicated.

The lens consists of a number of disc-shaped pieces of extremely transparent optical glass, called *elements*, with rounded surfaces that have been carefully molded, ground, and polished to microscopic tolerances. These elements are mounted in a cylindrical barrel such that they are all aligned along a common axis perpendicular to the plane of each element. Most lenses contain from five to seventeen elements. Each surface (front and back) of each element may be flat (*planar*), curved outward (*convex*), or curved inward (*concave*). Convex and concave surfaces are usually congruent with the surface of a sphere, but some elements may have a more complex curvature (such as parabolic curvature); such elements are said to be *aspheric*.

The lens gathers light reflected from a scene to be photographed and projects that light onto the film in the form of a reduced, sharp image, such that the film is exposed in a pattern that accurately records the scene. The lens bends incoming light rays so that light from any one given point in the image arrives at one and only one point on the film. Were it not for this focusing function of the lens, light would simply arrive as a blur on the surface of the film, and the entire image on film would be a uniform gray, instead of a recognizable representation of the image that was photographed. The principle of a lens is exactly similar to what you observe when you focus a miniature image of a brightly-lit scene onto a piece of blank paper with a simple magnifying glass. Indeed, a magnifying glass is really nothing more than a single-element optical lens—but photographic lenses are of far better quality and much greater complexity than a simple magnifying glass, in order to provide sharper, clearer images.

Photographic lenses contain more than one element for a number of reasons. The first and most important reason is that a single-element lens cannot be perfect; such a lens will always show various distortions and aberrations in the image it projects, because that is a fundamental limitation of a glass lens (a single glass element cannot focus all colors of light equally, for example). By using multiple elements, however, the defects of one element can be largely cancelled out by complementary defects in another element, such that the overall image is virtually free of defects. Other reasons for multiple elements include the need to adjust the focus distance of the lens, and the need to change the focal length of the lens (both of these concepts are explained further below).

Optical Quality

The lens characteristic that is most important to the quality of the final photograph image is its fundamental optical quality. Optical quality is a function of many parameters, including *resolution*, *contrast*, *distortion*, *flare*, *aberration*, and *bokeh*. A thorough treatment of any of these would fill several volumes and is thus beyond the scope of this paper, but we will attempt to summarize the importance of each of these different parameters here.

Resolution is the ability of a lens to focus fine details sharply onto the film. The simplest way to express lens resolution is in terms of the number of fine lines that the lens can sharply focus onto the film. The lines are assumed to be in pairs: one black line and one white line in each pair. The resolution is thus given in *line pairs per millimeter* (abbreviated as *lp/mm*) on film. A lens that resolves 50 lp/mm is able to sharply focus 50 pairs of white and black lines into each linear millimeter of the film. For a 35mm film frame, this represents a total of 1800 line pairs across the entire width of the frame.

Virtually no lenses are capable of maintaining the same resolution throughout the film frame. In most cases, resolution is much better in the center of the image than at the edges, and the corners are usually the worst of all. A lens that provides “flat” resolution (the same resolution throughout the image frame) tends to be very expensive, as such a lens is time-consuming to design and requires careful manufacturing techniques.

Contrast is the ability of a lens to accurately project the intensity of light and dark areas of the real world onto the film. In the world of photographic lenses, the higher the contrast, the better. An ideal lens would allow infinite contrast; that is, it would be able to accurately project blindingly-bright image areas right next to pitch-black image areas on the film, without any tendency for the bright areas to dim, and without any tendency for the dark areas to show any sign of "leaking" light from the bright areas. No lens has infinite contrast, but very good lenses tend to have very high contrast. Cheap lenses provide poor contrast, and as a result images photographed with them tend to look flat and cloudy, with little difference between light and dark areas.

Distortion is the tendency of a lens to distort the shape of the image as it is projected onto the film. For example, an inexpensive lens may tend to project straight lines in the real-world scene being photographed onto the film as curved lines. This is especially obvious with wide-angle lenses, in which it is difficult to completely avoid distortion. The better the lens, the less distortion there will be in the image it projects.

Flare is the tendency of internal reflections of light within the lens to cloud the projected image. Flare usually appears as large circles of diffuse light arranged diagonally across the image, often with the same shape as that formed by the iris diaphragm in the lens that controls the aperture. Flare is normally a bad thing in a lens, but some photographers in some situations consider it aesthetically interesting, especially for photographs taken with the camera pointed nearly towards the sun. Indeed, some digital image-manipulation programs used for retouching photographs actually provide a feature to simulate flare for precisely this reason.

Despite this, good lenses are always designed to reduce flare as much as possible.

Aberration is a class of characteristic defects in the projected image resulting from poor lens quality or design. The most common types of aberration are *spherical* aberration, appearing as a sort of smearing of image details towards the corners of the image; and *chromatic* aberration, appearing as color fringes around image details (particularly bright details). Spherical aberration is caused by the design of a lens; it is almost impossible to eliminate completely. Chromatic aberration is caused by the fact that optical lenses bend light of different colors to different degrees, which prevents them from all focusing at exactly the same point on the film—it is most obvious in telephoto lenses.

Spherical and chromatic aberration can both be enormously reduced by proper lens design and manufacture, but lenses with very low aberration are extremely expensive.

Bokeh (or *boke*, from a Japanese word meaning "blur" or "out of shape") is a term coined to describe the way out-of-focus areas look when projected by a lens. The details of the image area at the focus point of a lens are naturally projected onto the film as tiny points of light by any good lens; however, out-of-focus details are projected onto the film as overlapping circles of light, and together they combine to produce a blurred version of those details; Figure 3 is an example of a portion of a photograph that is out of focus and shows this phenomenon. The shape and character of these circles of light vary from one lens to another, which affects the aesthetic impression created by the blurred portion of the image—and this impression is what bokeh describes. For portraits and certain other types of photographic work, the appearance of out-of-focus areas outside the subject of the image may be artistically important, and so the bokeh of a given lens can be a significant factor in the choice of a lens for some photographers.

Figure 3

Focusing Adjustment

One of the fundamental limitations of any optical lens is that it cannot focus everything in a distant scene onto the film at once. At any given time, the lens can only focus portions of the scene that are at one specific distance from the lens; this distance is called the *focus point* (other terms are sometimes used, but this is the term we will use throughout this paper). Other portions of the scene in front of or behind this critical focus point are blurred, and the farther they are from the focus point, the more blurred they appear. See Figure 4 for an illustration of this. In the figure, notice that the only tree on the left that is truly in focus is the tree that is at the focus point; the other trees, both in front of and behind the focus point, are blurred in the image projected onto the film.

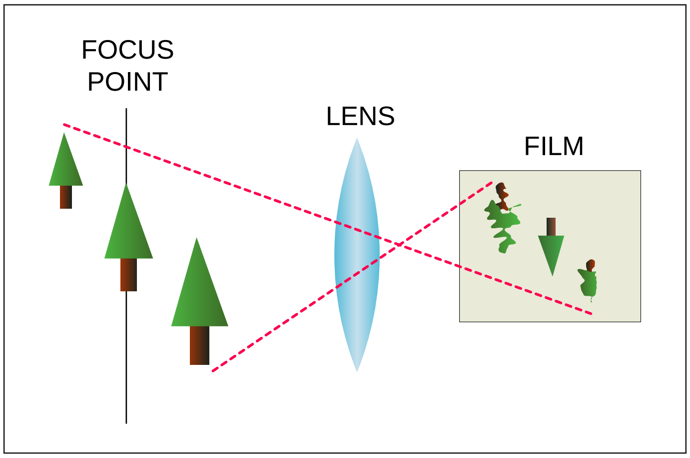

Figure 4

Because a lens with a fixed focus point would be very inconvenient to use (you could only use it to photograph subjects at one specific distance from the camera), photographic lenses normally contain one or more elements that can be moved within the lens assembly to change the focus point. A ring around the outside of the lens can be turned to move these elements and adjust the focus point for a specific photograph. Turning the focus ring moves elements forwards or backwards via a mechanical, screw-like arrangement (a *helicoidal thread*) inside the lens barrel, and this change in the relative positions of elements within the lens changes the focus point.

Some cameras incorporate computer-driven mechanisms that can adjust the focus of the lens automatically to match the distance of the subject in a photograph; lenses containing such mechanisms, or compatible with similar mechanisms within the camera body (which operate through a simple mechanical linkage with the focus ring of the lens) are called *autofocus* lenses; and cameras that provide this feature are said to be equipped with autofocus (AF).

All lenses can be adjusted to focus at “infinity,” meaning that the focus point is set to be infinitely distant from the camera. In practice, nothing is ever infinitely far away, of course, but scenes such as distant mountains, stars, and even buildings on a skyline are usually “infinitely” distant from the point of view of most lenses, so the infinity setting is normally used for these types of photographic subjects.

Depth of Field

Although a lens can only focus for one specific distance at a time with complete precision, in most cases objects that are very close to the focus point in the scene being photographed will be only slightly blurred. In fact, within a certain range of distances in front of and behind the focus point, the blurring is so slight that it is not visible to the unaided human eye. This range of distances within which objects appear to be in sharp focus (at least to the human eye) is called the *depth of field*; as long as a part of the scene is within the range of distances covered by the depth of field, it will appear to be in correct focus (even if it is not exactly at the focus point). The depth of field in any given situation is a function of several factors, including the actual focus-point distance (depth of field is greater when the focus is on very distant objects), the current aperture of the lens (the smaller the aperture setting, the greater the depth of field will be), and the focal length of the lens (the shorter the focal length, the greater the depth of field—see below).

Some lenses, particularly lenses with manual focus adjustments only (no autofocus), are equipped with depth-of-field indicators near the focusing ring on the lens. These indicators show the range of distances that will be approximately in focus for any given setting of the focus point. Autofocus and zoom lenses may omit the depth-of-field indicators because they are rarely used and can be rather complex to implement on these lenses (zoom lenses are explained later in this paper). Additionally, in some cameras, it is possible to verify the depth of field simply by looking through the camera’s viewfinder at the scene to be photographed (more on this later).

All lenses have a focus point called the *hyperfocal distance*. This focus point is the distance that gives the greatest depth-of-field overall; when the lens is focused to the hyperfocal distance, everything from half the hyperfocal distance to infinity will be in apparent focus. The hyperfocal distance is not explicitly indicated on the lens, but it can be inferred from the depth-of-field and focus indicators on the lens barrel.

Focal Length

Every lens has a built-in characteristic called *focal length*. Technically, the focal length is the distance behind the lens at which the lens will focus a sharp image of an object infinitely distant from the front of the lens (such as a distant star in the sky, which can be considered infinitely distant for optical purposes), when the focus point is set to infinity. A 50-millimeter lens, then, is a lens that must be positioned 50 millimeters in front of the film surface in order to correctly focus a scene that is infinitely far away.

For the sake of simplicity in this paper, it is sufficient to remember that the focal length is an indication of the field of view of the lens—that is, the width of the image that the lens can focus on the film when the subject is at a given distance from the camera.

For example, if you photograph, say, a group of trees fifty feet away with a lens that has a short focal length, all of the trees may appear in the image, and they will appear relatively small compared to the image overall. In contrast, if you photograph the same group of trees from the same spot using a lens with a long focal length, you may only see part of the trunk of a single tree in the image, and that small portion of the one tree will fill the entire frame of the image. See Figure 5 for an illustration.

In other words, long focal lengths tend to "magnify" an image, like a telescope (in fact, telescopes magnify distant images precisely because they contain lenses with very long focal lengths). Short focal lengths tend to make objects in an image look small and distant—the opposite effect from that of a telescope. Additionally, long focal lengths tend to flatten apparent perspective, so that objects at different distances appear to be of the same size, whereas short focal lengths tend to exaggerate perspective, so that even objects at only slightly different distances from the camera seem to be of dramatically different sizes.

In the most common type of camera, a lens having a length of 50 millimeters is said to be "normal," meaning that photographs taken with this lens, when seen from a fairly typical viewing distance, will show approximately the same perspective and field of view as what one would see in real life from the position of the camera. Lenses with focal lengths shorter than 50 millimeters are often called *wide-angle* lenses. Lenses with long focal lengths are typically called *telephoto* lenses.

An example of the effects of focal length is illustrated in Figure 6 on the next page. The top photograph in the figure was taken using a lens having a focal length of 35 millimeters. The bottom photograph of the figure was taken with a lens having a focal length of 135 millimeters. Both photographs were taken from exactly the same spot. Notice that the perspective does not actually change with the focal length of the lens, but the portion of the scene recorded on the film changes considerably, which creates the illusion of a change in perspective (mainly in the form of a "flattening" of perspective as the focal length increases). The wide-angle 35-millimeter lens also shows quite a bit of the scene at reduced size (the bridge and pier are visible, and the boat in the center, while visible, is very small), whereas the telephoto 135-millimeter lens shows only the central portion of the scene, greatly magnified (the boat and bridge behind it are large and nearly fill the frame,

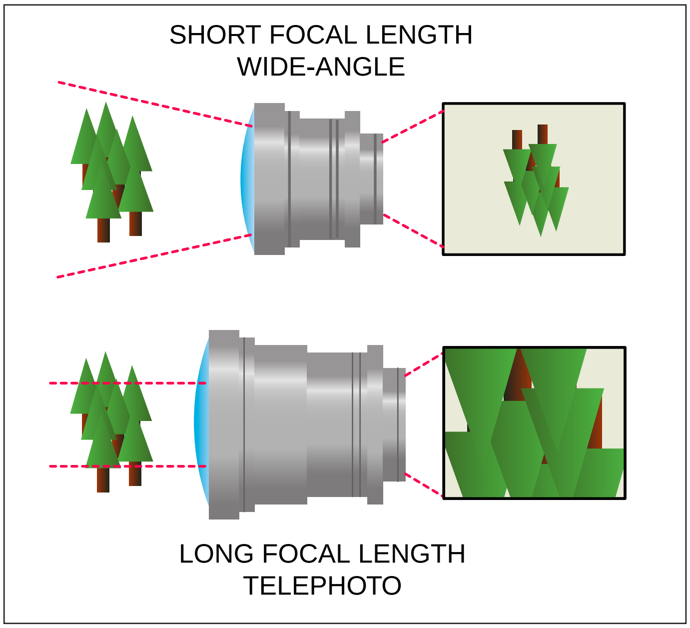

Figure 5

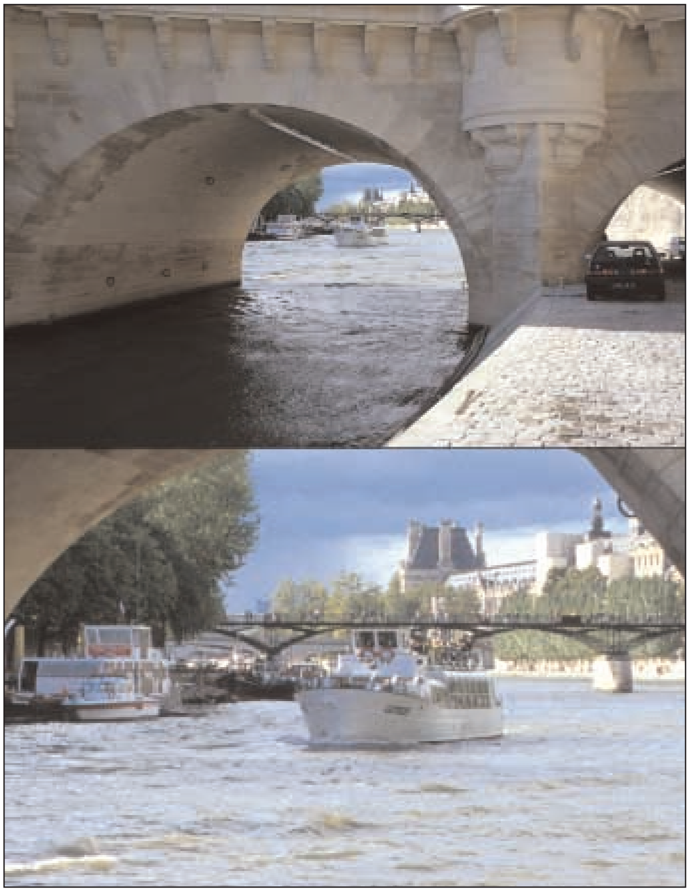

Figure 6

whereas the bridge and pier in the foreground with the 35-millimeter lens are no longer even in the frame at 135 millimeters).

In theory, the physical length of a lens is closely related to its focal length. However, this can make very long focal lengths and very short focal lengths impractical to construct, since the former become so huge that they are hard to transport and support, and the latter become so flat that they can barely be mounted on the camera. Because of this, modern lenses incorporate special design features that allow long focal lengths to be “telescoped” into a shorter physical package (whence the name *telephoto* for lenses of long focal lengths), and others that allow short focal lengths to be “stretched” a bit to make them easier to mount on the camera. In practice, then, lenses tend to be more consistent in physical size than their variable focal lengths might theoretically imply.

Prime and Zoom Lenses

Traditionally, a lens has a single, fixed focal length. Such a lens provides optimal image quality for the focal length incorporated into its design. However, many of today’s lenses have a variable focal length; that is, they can be adjusted to any focal length between an upper and lower limit defined by the design of the lens. Such a lens is called a *zoom* lens (as opposed to a standard or *prime* lens, which is always of a single, fixed focal length). Thus, a zoom lens designated as 28-70 in focal length can be set to any focal length between 28 and 70 millimeters. The focal length is manually adjusted using a ring on the lens barrel, in the same way that focus is adjusted. Some zoom lenses are “one-touch” lenses that use the same ring for focus and zoom: turning the ring changes the focus, and pushing or pulling it along the length of the barrel adjusts the focal length. Unlike focus adjustments, there is no such thing as an automatic zoom in photography; the focal length is always adjusted manually by the photographer as she sees fit.

Zoom lenses have the obvious advantage of convenience: a single zoom lens with a range of 28-70 millimeters, for example, can replace three prime lenses of 28, 50, and 70 millimeters (as well as an infinite number of prime lenses with other intermediate focal lengths). Despite this, however, prime lenses remain popular, mainly because primes generally provide better image quality than zooms. A zoom lens usually requires more complex and delicate construction than a prime lens, and it is very difficult and expensive (albeit theoretically possible) to construct a zoom lens that provides the same optical quality as a prime lens throughout its focal-length range. Modern-day zooms closely approach the quality of primes in many cases, but for pure image quality alone, primes still have the advantage, and so some photographers avoid zoom lenses in favor of prime lenses, despite the inconvenience of having to physically change lenses in order to obtain different focal lengths.

Since there is a relationship between the physical length of a lens and its focal length, zoom lenses often change in physical length as the focal length is adjusted, with longer focal lengths increasing the physical length of the lens, and shorter focal lengths decreasing the physical length. Advanced lens designs limit the change in length to a fraction of what might be theoretically expected for the corresponding change in focal lengths, but in some inexpensive lenses and some very wide-range zooms, the lens may double or triple in length from one end of the focal-length range to another.

Diaphragm, Aperture and Speed

As previously mentioned earlier in this paper, the lens contains an aperture or opening that restricts the amount of light entering the lens that is allowed to reach the film. The aperture of a lens is usually adjustable, and normally takes the form of an *iris diaphragm*, which is a set of thin, curved, overlapping metal blades in a circular arrangement that can be moved inward or outward with respect to the center of the lens. Their overlapping pattern produces a roughly circular hole (the aperture) surrounded by darkness that looks schematically very much like the iris of a human eye (the iris being the colored part of the eye that surrounds the dark pupil in the middle). Figure 7 shows how the diaphragm appears within the lens for two different aperture sizes (the numbers appearing beneath each illustration are explained under *F-stops*, below).

The adjustment of the aperture of a lens is accomplished manually by turning a ring on the lens, but most modern-day lenses also allow for automatic adjustment of the aperture by the camera body. This latter feature makes it possible for a computer within the camera body to automatically adjust aperture and shutter speed for a given situation to optimize exposure of the film. In addition, most cameras that allow the photographer to view the scene to be photographed directly through the lens are designed to hold the aperture wide open at all times except during the instant of exposure, in order to make the image in the camera viewfinder brighter, thereby facilitating composition.

In essence, aperture is simply a measure of the amount of light reaching the film through the lens. The smaller the aperture, the less light is admitted by the lens, and the lower the exposure of the film to light at any given shutter speed. All lenses have a maximum and a minimum aperture, determined by the limits of the mechanical movements of their diaphragms and by the design and placement of their optical elements. Small apertures are relatively easy to provide from a lens-engineering standpoint, and so almost all lenses provide a fairly small minimum aperture; however, large apertures present many problems in lens design, and so the maximum aperture of a lens often is closely linked to the quality of its design, and this in turn is often closely linked to the cost of the lens. Lenses with very large maximum apertures are extremely expensive to design and construct. They also tend to be very large in diameter, since the diameter of the aperture must be close to the focal length for large apertures (that is, lenses with very large maximum apertures are almost as wide as they are long).

Because a larger aperture allows film to be exposed properly at higher shutter speeds (shorter periods of exposure), a lens with a large maximum aperture is said to be “fast,” and a lens with a small maximum aperture is said to be “slow.” Following the same logic, the maximum aperture of a lens is conventionally referred to as its *speed*. The minimum aperture of a lens is usually more than adequate for any purpose, and so it is not normally specified when characterizing a lens.

Zoom lenses occasionally have variable maximum apertures; that is, the maximum aperture does not remain the same throughout the zoom range (typically, the maximum aperture or “speed” of the lens diminishes as the focal length is increased). This is the case because it is difficult and expensive to design a zoom lens that has a constant maximum aperture throughout its focal-length range, and a constant maximum aperture is rarely a critical requirement, especially for novice and amateur photographers. Professional (and expensive) zoom lenses often provide a constant maximum aperture throughout the focal-length range, however.

F-stops

Aperture in a lens is designated by a number that

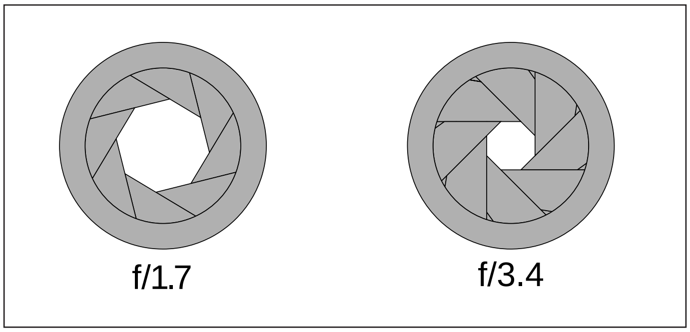

Figure 7

represents the ratio between the diameter of the aperture (which is a roughly circular hole) and the focal length of the lens. This number is called an *f-stop*, and it is usually written *f/n*, where *n* is the f-stop number. For example, f/2 (pronounced "eff-two") indicates that the diameter of the aperture in the lens is equal to the focal length divided by two. The reason this system of designating aperture is used is that the real aperture of a lens varies with the focal length (a one-inch aperture in a lens with a focal length of 50 millimeters allows the same amount of light to reach the film as a two-inch aperture in a lens with a focal length of 100 millimeters), which makes directly specifying the diameter of the lens aperture awkward (because the focal length must be specified as well). By specifying the f-stop as a fraction incorporating the focal length, however, one can always be sure that an f-stop of f/2 represents a constant amount of light reaching the film, no matter what the actual diameter of the aperture or the focal length of the lens might be.

F-stops are marked on the aperture ring of a lens as a succession of numbers representing the focal length divided by the aperture diameter. The numbers are always the same, and Table 1 shows the sequence in which they appear.

The reason for the seemingly unusual sequence of f-stops on the aperture ring is that each corresponds to exactly half the exposure of its predecessor, and twice the exposure of its successor. For example, when the aperture is changed from f/1.4 to f/2 (a move forward of one f-stop), the amount of light entering the lens is cut in half; and when the aperture is changed from f/16 to f/11 (also a change of one stop), the amount of light entering the lens is doubled. Each additional f-stop further multiplies or divides the exposure by two. Since the f-stop represents the ratio between the focal length and the *diameter* of the aperture, and since the exposure depends on the *area* of the aperture, the f-stop is the square root of the area (because area varies as the square of diameter), which is why it consists of numbers that do not seem intuitively meaningful (except to math geeks). This is probably not the ideal way to mark lens apertures, but it has been in use for such a long time that it is unlikely to change.

Standard F-Stop Numbers			
1	2.8	8	22
1.4	4	11	32
2	5.6	16	45

Table 1

This use of distinct f-stop settings on aperture rings has led to the use of the term *stop* in photography to generically designate any change in light or exposure by a factor of two. In other words, increasing exposure by *one* stop means *doubling* exposure (multiplying it by two); decreasing exposure by *three* stops means dividing it by *eight* (2x2x2); and so on. Similarly, it is customary to say that one part of an image or one lighting situation is a certain number of stops brighter or darker than some other. The total range of light intensity encountered in everyday life, from starlight to the surface of the midday sun, represents a ratio of about ten million to one, which corresponds to about twenty stops. Excluding direct shots of the sun, most photographic situations fit into a range of about fifteen stops. In order to be successfully exposed on film, however, the total range of light levels in a scene should not exceed five to seven stops, as this is the maximum range that film can record in a single exposure while still showing a useful level of detail in shadow and highlight areas.

Miscellaneous Lens Characteristics

The major lens specifications of importance to photographers have already been described: optical quality, speed, and focal length. There are, however, some lesser characteristics and types of lenses that can be important in specific situations.

Macro lenses are intended specifically for photographing subjects at extremely close range (a few centimeters or inches). The design of such lenses is optimized to permit and enhance use at such close range. Some normal lenses provide a "macro" position on the zoom or focus ring that allows their use as macro lenses, but their image quality may not be quite as high as that of a purpose-built macro lens. Most ordinary lenses are unable to focus on objects closer than half a meter (20 inches) or so.

Reflex telephoto lenses are lenses of extremely long focal length that incorporate some of the principles of telescope optics in order to reduce the diameter and especially the length of the lens. Reflex telephoto lens designs are often seen for lenses of more than 600 mm in focal length, and they use a combination of traditional lens elements and mirrored surfaces to compress the optical path lengthwise and allow construction of lenses in very long focal lengths with reasonably short barrels (although they tend to be quite wide).

Retrofocus wide-angle lenses are lenses of extremely short focal length that are specially designed to make them a bit longer than they would otherwise be. This is necessary because lenses with very short focal lengths must be so close to the film that they cannot be easily mounted on the standard lens mount of the camera, and they may even interfere with the movement of parts inside the camera body. By making certain adjustments in the optical design of such lenses, they can be lengthened in a way that allows them to be held at a greater distance from the film without increasing the effective focal length.

Fisheye lenses are lenses of very short focal length (eight to ten millimeters, usually) that look like giant fish eyes when mounted on the camera (whence their names). They can fit 180° of the image field onto the film (from one horizon to the other, for example), but they distort the image dramatically in a characteristic way. They are used for special-effect photographs and for photographs in which it is essential that an extremely wide field be included in the image, even with distortion (some types of surveillance photography, for example).

Aspheric lenses, which have been previously mentioned, incorporate lens elements with surfaces that do not fit on a sphere. Aspheric lenses are able to avoid some of the optical defects that are a problem with lenses using purely spherical elements. In some cases, aspheric lenses are used to save money (as in small, compact cameras); in other cases, they are used to improve the optical performance of high-end lenses. As a result, the mere fact that a lens is aspheric does not necessarily mean that it is optically or mechanically better or worse than any other lens. The method used to produce the aspheric elements may be an indication of quality, however: the cheapest aspheric elements are molded from transparent plastic, whereas the most expensive are painstakingly ground from optical glass using special equipment that can vary the curvature of the ground surface as required (at one time, they had to be ground entirely by hand, but fortunately that is no longer necessary).

Figure 8

Perspective-control (PC) lenses are lenses that can be adjusted such that the barrel of the lens is no longer perpendicular to the film. To understand the utility of such a lens, one must remember that the perspective of an image in a photograph depends exclusively on the position of the film plane with respect to the image being photographed. When a large building or some similar subject is photographed from a spot very near the subject and on the ground (as is likely if the photographer is standing in front of a building with the camera in his hand or on a tripod), strong perspective will cause the top of the building to seem much smaller than the bottom. This occurs because the film plane in the camera is not parallel to the front of the building, and the film plane is not parallel because the photographer must tilt the camera in order to get the entire facade of the building into the photograph. A PC lens corrects for this by allowing the photographer to aim with just the lens, while keeping the camera body (and thus the film plane) parallel to the front of the building. The net result is that the strong perspective in the image is eliminated.

Because architectural photography is virtually the only type of photography in which strong perspective is a real problem, the use of PC lenses is pretty much restricted to architectural work. See Figure 8 on the previous page for an example of a photograph taken with an ordinary lens (top image) and a PC lens (bottom image).

Apochromatic lenses are lenses that incorporate special correction for *chromatic aberration*, which is a type of lens defect—present in all lenses to some degree—that causes light of different colors to be focused unequally on the film. Truly apochromatic lenses are difficult and expensive to design and build, and they are usually used only for long lenses (such as telephoto lenses), for which chromatic aberrations tend to be much more of a problem.

The Camera Body

Another essential part of a camera is the camera body. This part of the camera is really just a box that contains the film and protects it from light (except during exposure). In practice, however, modern camera bodies have become extremely complex, often containing computers, motors, sensors, light meters, liquid-crystal (LCD) displays, and much more. In addition to protecting the film and controlling its exposure, the camera body often functions as a command center for the camera as a whole, by directing automatic adjustment of lens focus and aperture, by triggering external flash units, and so on.

Rangefinders and SLRs

In 35mm photography (so called because of the width of the film strip used by the cameras), which is by far the most common type of film photography in use today, two types of camera bodies predominate: the *rangefinder* body, and the *single-lens reflex* or *SLR* body. The main difference between the two types of bodies is in the way the photographer views the image he is photographing (see Figure 9). In a rangefinder, the photographer looks through a small window in the camera body that shows him the approximate composition of the image that will be focused onto the film (F in the figure) when the shutter is triggered. In a single-lens reflex, the photographer looks at a tiny screen showing a direct projection of the image to be photographed as it is actually seen by the lens, reflected from the lens path by a mirror (M in the figure); the mirror moves out of the way briefly when a photograph is taken in order to allow light to strike the film behind it. Both types of camera body have their respective advantages and disadvantages.

Rangefinders are so called because the focusing mechanism in such cameras involves the use of an optical rangefinder, a device that measures distance by measuring parallax between images seen from two different points. For the photographer, focusing a rangefinder involves visually aligning two overlap-

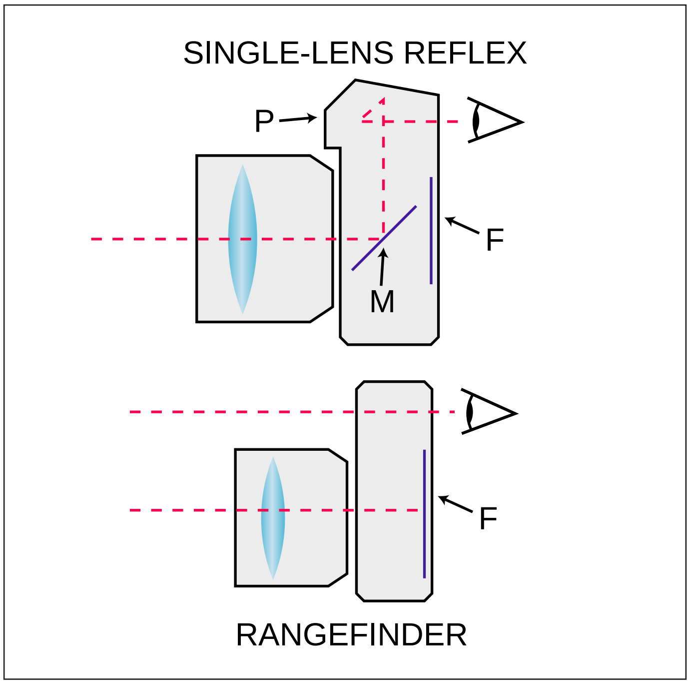

Figure 9

Figure 10

ping versions of the image to be photographed until they coincide in the viewfinder; the photographer turns the focus ring on the lens until this alignment appears correct. This type of focusing is more accurate than that used by SLRs for lenses of focal lengths up to about 90 mm or so.

A typical professional SLR (a Nikon F5) is pictured in Figure 1, at the beginning of this paper. Figure 10 shows a typical professional rangefinder (a Leica M6 TTL).

By extension, any camera that uses a small viewfinder independent of the main lens for composing the photograph tends to be called a rangefinder, even if it does not contain an actual rangefinder mechanism for focusing. Many compact amateur cameras resemble rangefinders in their use of a separate viewfinder, but they do not actually contain a rangefinder mechanism.

Single-lens reflex cameras are so called because both the viewfinder and the film see the world through the same, main photographic lens (thus a *single lens*), and because the image captured by the lens is bounced around (*reflexed*) on the way to the viewfinder so that the photographer can see the image in its real-world orientation (that is, not backwards or upside down, as it is projected onto the film). In the SLR, when film is not actually being exposed, the image from the lens is reflected onto a small screen in the viewfinder by a mirror, and then via a number of intermediate mirrors to correct the orientation of the image. The photographer can thus see the actual image that will be recorded onto film, instead of the approximate image shown by a rangefinder. The mirror arrangement is usually contained in a structure called a *pentaprism* (P in Figure 9) that is mounted over the lens and is responsible for the characteristic “hump” of SLR cameras. Focusing an SLR require simply turning the focus ring on the lens until the image in the viewfinder looks sharp and clear. This type of focusing is more accurate than rangefinder focusing for lenses of long focal lengths (greater than 90 mm).

Both rangefinders and SLRs produce images of identical quality, so the choice between them is usually purely a matter of photographer preference. There are numerous differences between the two types of camera from the photographer’s standpoint that can influence her preference. For example, rangefinder cameras tend to be very small, silent, and discreet, whereas SLRs tend to offer much higher shutter speeds. Rangefinders allow the subject to be viewed even during exposure (because the viewfinder is independent of the lens used to expose the film, and there is no mirror to depend upon), whereas SLRs do not (because the photographer’s view through the lens via the mirror is temporarily interrupted while light is diverted to the film for exposure). Still another difference is that rangefinders offer only an approximation of the actual image frame being recorded on film, whereas SLRs show the photographer exactly what is being recorded.

Both types of camera exist in all price ranges; however, most inexpensive compact consumer cameras and all disposable cameras are of the rangefinder type (minus the actual rangefinder mechanism). Most professional 35mm cameras are of the SLR type, although a few of the finest professional 35mm cameras are also true rangefinder cameras.

The Shutter

The most important component of the camera body is the shutter. This is the mechanism that prevents light from the lens from reaching the film until the moment of exposure. During exposure, the shutter briefly opens to allow light to reach the film, and then closes again when exposure is concluded. Modern shutters operate with extraordinary speed, often opening and closing within only a few thousandths of a second.

The shutter of most cameras is adjustable: the photographer can select the amount of time that the shutter is to remain open during an exposure by turning a dial or pressing buttons. Like the aperture ring on a lens, the dial that controls shutter speeds typically shows increments in powers of two. For example, most shutter dials include speeds of 1/2 second, 1/4 second, 1/8 second, and so on, often down to 1/2000 second or less. To save space, the shutter-speed dial often shows only the lower part of the fraction; that is, the dial will show 2, 4, and 8 for 1/2, 1/4, and 1/8 second, respectively.

Virtually all 35mm camera bodies today incorporate *focal-plane* shutters. A focal-plane shutter is so called because it is positioned right in front of the film, and thus very near the *focal plane* (the plane of the film, at which the image to be photographed is focused by the lens). Focal-plane shutters typically consist of two curtains, each of which is made of opaque and flexible cloth, or of overlapping metal or plastic blinds or blades. Under normal conditions, one shutter curtain or one set of blades blocks the light from the lens and prevents it from reaching the film, and the other curtain or blade set remains out of the way to one side of the film. When the shutter is triggered to take a photograph, the curtain or blade set blocking the film moves aside, exposing the film to light from the lens. When the exposure is finished, the second curtain or blade set moves into place over the film, once again preventing any light from the lens from reaching it. When the camera mechanism is reset for the next exposure (by turning a lever manually in some cameras, or by an automatic motor in others), the curtains return to their original positions in such a way that the film is never exposed to light as they move. The entire process of exposure takes place in a small fraction of a second (although the winding of the mechanism for the next exposure may take a second or so). The curtains or blades must move at extremely high speed, and so they are typically light, delicate, and fragile. In some cases, they are even made of titanium, in order to maximize strength while minimizing size and weight.

Other types of shutters are used in very inexpensive and disposable cameras and a handful of professional cameras; the most common is a diaphragm-like arrangement that operates like the aperture adjustment in the lens, except that it has only two positions—open and closed—and that it is optimized for speed, not accurate positioning. In any case, whatever the type of shutter, the purpose is always the same: the shutter hides the film until the moment of exposure, and then briefly exposes the film to light from the lens for a precisely-controlled period.

Synchronization Speed

All focal-plane shutters have a *synchronization speed (sync speed)*. This is the highest speed at which the film is entirely exposed over all of its surface simultaneously.

At speeds *below* the sync speed, the shutter has time to open completely before it begins to close again, which means that, at some point during the exposure, the entire film surface is uncovered simultaneously. In contrast, at speeds *above* the sync speed (faster than the sync speed, in other words), the shutter never has time to uncover the entire film frame at once before it must start closing again.

For example, in the case of a focal-plane shutter, at speeds above the sync speed, the second curtain or blade set of the shutter will start to cover the film before the first curtain has completely finished opening, which means that the film is never more than partially uncovered; any given point on the film will still receive the correct exposure, but different points on the film will be exposed at different times, and at no point will the *entire* film frame be exposed to light at once. In other words, the shutter at speeds above the sync speed behaves much like a moving slit of variable width that travels across the film, exposing different parts of the film as it moves. The higher the shutter speed is above the sync speed, the smaller the slit will be. At shutter speeds below the sync speed, the “slit” is so wide that the entire film will be uncovered at once at some point during the exposure.

Why is this important? It usually isn’t, unless a *flash unit* is being used to provide additional light for the photograph. A flash unit is simply a very bright light that flashes on for a fraction of a second during an exposure to provide extra light for taking a photograph. The reason sync speed is important for the use of a flash unit is that it is essential that the film be completely uncovered *at the instant that the flash unit fires*—otherwise the light from the flash will not reach the part of the film that was not uncovered at that instant, and so part of the resulting image will be dark. The sync speed of a shutter is the maximum speed at which the film will be completely uncovered at some point, and so it is the maximum speed at which a flash unit can be used.

In general, the higher the sync speed of the shutter, the better a camera is (and the more money it costs). Some types of focal-plane shutter are capable of higher sync speeds than others; for example, shutters that use fabric curtains tend to have slower sync speeds than shutters that use overlapping blades, and of those shutters that use blades, those that use vertically-moving blades have higher sync speeds than those that use horizontally-moving blades.

Some other types of shutters have similar problems with synchronization speeds for flash, but the problem is most obvious with focal-plane shutters, in part because focal-plane shutters are so widely used.

Automatic Exposure Control

One of the most important features of modern camera bodies is some form of automatic exposure control. For any given subject that a photographer might wish to photograph, there is some combination of aperture size (in the lens) and shutter speed (in the camera body) that will expose the film in a way that will provide the best possible image when the film is developed. This combination of aperture size and shutter speed is called *exposure value* (EV) and must be calculated for every photograph taken.

Most cameras, except for inexpensive consumer models and disposable cameras, allow the photographer to manually adjust exposure value, by controlling the shutter speed and the lens aperture. (In fact, a few cameras, particularly certain professional models, *require* that the photographer adjust exposure manually.) The photographer adjusts each of these independently based on the exposure value that she has calculated to be appropriate for the image to be photographed. Photographers who work in this way use an instrument called an *exposure meter* or *light meter* to measure the amount of light in the scene to be photographed, and then calculate the proper exposure value (manually or with the aid of a calculator built into the meter). Most camera bodies that allow for manual adjustment of exposure include a convenient, built-in meter for this purpose that measures light directly as it enters the lens; others have no meter, and the photographer must depend on her own judgement, or use a separate meter.

In many situations, there is no advantage to calculating exposure values manually or with an exposure meter; the photographer's main concern is simply to get a proper exposure of the image on film. For these situations, modern camera bodies typically incorporate not only an exposure meter, but also a computer that reads the information provided by the meter and directly adjusts aperture and shutter speed to provide a correct exposure. The photographer simply aims the camera, and when she presses the shutter button, the computer in the camera body quickly calculates the correct exposure, based on the luminosity (brightness) of the subject at which she is aiming the camera, and then adjusts the aperture and shutter speed almost instantaneously. The camera then opens the shutter and makes the exposure. The entire automatic process takes only a small fraction of a second.

The mechanics of automatic exposure systems are complex, and their reliability varies enormously from one camera model to another. Top professional camera bodies are typically extraordinarily accurate in calculating exposure even in difficult situations, whereas inexpensive compact consumer cameras are only moderately accurate, and only in the simplest circumstances. Disposable cameras have no exposure control at all; they simply expose every image identically, using parameters that are likely to provide acceptable results in an average situation (such as exterior photographs taken in bright sunlight).

In camera bodies with automatic exposure capability, there is often a way for the photographer to turn off the automatic features in whole or in part if she wishes to take control of certain parameters herself. Thus, many cameras include an *aperture-priority* operating mode that allows the photographer to select an aperture, while the camera body calculates the correct shutter speed in consequence (based on information from its built-in exposure meter). The main advantage of aperture-priority is that it allows the photographer to control depth of field without having to worry about making the corresponding adjustments in shutter speed (the latter being made by the camera automatically). For example, in any given situation, if the photographer wishes to reduce the aperture by half in order to increase depth of field, the shutter speed must be halved also in order to expose the film for twice as much time and compensate for the fact that only half as much light is reaching the film through the lens; aperture priority allows the photographer to worry only about the aperture, letting the camera's computer set the shutter speed in consequence.

Some cameras also provide a *shutter-priority* mode in which the photographer chooses a shutter speed and the camera body selects the corresponding aperture. A shutter-priority mode is not as common as an aperture-priority mode because there are fewer reasons for a photographer to want to control shutter speed directly.

In consumer compact cameras, there is often no way to override automatic exposure—everything is controlled by the camera's computer. In more advanced and expensive cameras, though, it is possible to completely shut off automatic exposure so that the photographer has full manual control over shutter speed and aperture. Of course, once the automatic features are shut off completely, it is up to the photographer to calculate the correct exposure parameters—the meter with the camera still operates, but it's up to the photographer to use it and adjust exposure appropriately. Some photographers prefer to leave everything on automatic, others prefer to do everything manually, and many simply vary their practices as a function of the shooting situation.

Automatic exposure control in the camera body requires that some sort of linkage exist between the camera body and the lens, so that the camera body can adjust the lens aperture itself. Linkages like this have existed for many years on most 35mm cameras and it can usually be taken for granted that any lens that can be mounted on a camera with automatic exposure control necessarily contains the proper linkage for aperture adjustment.

Some of the most recent cameras involve a much more elaborate communication between lens and camera body. In these cameras, the lens can tell the camera body the distance of the subject (based on the focus distance set by the photographer or the camera), the current zoom setting, if applicable, and the current maximum aperture. This allows the camera body to be even more accurate in its calculation of exposure values, particularly when a flash unit is used to add additional light for exposure.

Autofocus

Autofocus—the automatic adjustment of the focus distance in the lens of a camera—has been mentioned previously in this paper in the discussion of lenses. Autofocus is an extremely common feature among modern 35mm cameras. The focus adjustment takes place in the lens, since that is where the optical elements that control the focus point are located; however, this adjustment is driven by the camera body, which calculates the correct focus distance by actually examining the image to be photographed and then commands the adjustment of the lens focus.

Since the camera must adjust the lens focus in an autofocus system, there must be a way to turn the focus ring automatically. The most common method for accomplishing this is to use a tiny motor in the camera body to drive a gear train that moves the focus ring via a shaft that is part of the lens mount. This works well for adjusting focus, but it is somewhat noisy and slow. In consequence, more expensive professional cameras and lenses often incorporate a focusing motor within the lens itself, and the camera body controls this motor remotely using electrical signals sent through contacts in the lens mount. In addition, these more expensive lenses often use special piezoelectric ring motors (sometimes called "ultrasonic" motors) that provide extremely precise, positive, rapid, and quiet operation. The fastest mechanisms of this kind can silently adjust a lens to correct focus within a fraction of a second.

When it comes to determining the correct focus, many techniques are used. Some cameras (particularly consumer compacts) bounce an invisible infrared light off the subject of the photo to measure the distance of the subject from the camera. A few cameras emit an ultrasonic beep and time its return to the camera to determine the subject distance—rather like the sonar of a submarine. Still other cameras measure the contrast in the viewfinder, and adjust the focus until the contrast is maximal (a blurry image shows only gradual transitions between light and dark, whereas a sharp image shows obvious, high-contrast transitions between details). Finally, the most advanced 35mm cameras with autofocus use a kind of electronic rangefinder system; this system operates much like the rangefinder focusing previously discussed, except that a computer in the camera body adjusts the focus while examining the overlapping images, instead of the photographer. The rangefinder mechanism provides the greatest accuracy and speed for autofocus; the best of these autofocus systems can continuously maintain the correct focus on a moving subject in real time.

Most autofocus mechanisms do not allow the photographer to override them without turning the autofocus feature off. However, professional lenses with piezoelectric drive motors usually allow the photographer to adjust the focus herself after the autofocus is completed (the latter takes only a fraction of a second), and some such lenses even allow the photographer to override the autofocus as it is operating.

On professional and high-end consumer camera bodies, autofocus mechanisms frequently offer a choice between a "one-time" mode and a "continuous" mode. In the former, which is usually the default mode, the camera focuses once when the shutter button is partially depressed, and then remains at that focus setting until the button is released and slightly depressed again. In the latter mode, the camera focuses continuously for as long as the shutter button is depressed, adjusting focus whenever the subject moves or changes. Some cameras allow the photographer to set the focus at a specific distance and arm the shutter, and then the camera triggers the shutter automatically whenever something moves into the plane of focus previously set. Finally, some cameras even provide a "follow-the-subject" mode, in which the camera follows the subject when it moves, adjusting focus as necessary, after the photographer has focused on it initially.

Because it is impossible for an entire image to be in focus simultaneously (unless everything in the image is at the same distance from the lens), it is necessary to tell an autofocus system which part of the image must be in focus. Most cameras with autofocus do this by defining one or more small zones in the viewfinder that the photographer can select as target zones. When the autofocus operates, it will attempt to focus on whatever happens to fall into the selected target zone, ignoring the rest of the image. There is always at least one zone, usually in the center of the image, and most cameras offer a choice of several zones at different points in the image frame. Another system, used primarily by certain cameras made by Canon, actually tracks the movements of the photographer's eye against the viewfinder and focuses the camera on whatever the photographer looks at in the image.

Program Modes

Many modern cameras offer *program modes* that simultaneously adjust aperture, shutter speed, and focusing mode to match specific shooting situations. For example, a camera might offer a "sports" mode for shooting sporting events, and a "portrait" mode for shooting portraits. The utility of these program modes is controversial. The most elaborate program modes are offered in mid-range cameras, with professional cameras offering a far more limited select of modes, and inexpensive cameras offering none at all.

The Film

Finally we come to the last major part of the camera: the film. The major difference between the film and the other parts of the camera (lens and camera body) is that film is not reusable, whereas the lens and camera body can be reused indefinitely.

Film is the medium upon which images are recorded by the camera, by the process previously summarized in this paper. Physically, film takes the form of a sheet, strip, or plate covered with a very complex, multilayered chemical coating called an *emulsion*. The emulsion is extremely sensitive to light, and undergoes a permanent change when exposed to it. The degree of change in the emulsion depends on the intensity of the light to which it is exposed and the duration of that exposure (*i.e.*, aperture and shutter speed). The effect of multiple exposures is cumulative, and normally a given area of film is exposed only once. By projecting a focused image onto the film, it is possible to record images permanently, because the bright areas will be more exposed than the dark areas, in exact correspondence with the bright and dark areas of the image that was photographed.

Film Formats

Film is manufactured in a number of formats. By far the most popular format for film is the 35mm format, which is why we have emphasized this format in this paper. Film in the 35mm format consists of a strip of flexible plastic (usually cellulose triacetate) that is 35 millimeters wide and of variable length, with small sprocket holes punched near both edges. The strip is wound into a disposable, light-proof cartridge that can be loaded into a camera body. A short leader section of film protrudes from the cartridge and allows the film to be wound onto a take-up spool inside the camera body (some camera bodies do this automatically, while others require that the photographer thread the leader into the take-up spool). Once the camera body is closed to protect the film, the film is unwound little by little, using a wheel that engages the sprockets. The film is drawn past the area containing the shutter, called the *film gate*, and it is held there as each image is exposed by projecting light onto the film from the lens. After each exposure, the film is wound so that the exposed frame of film moves away from the film gate and is replaced there by a fresh, unexposed section of the film. When the entire strip of film has been exposed in this way, it is wound back into the light-proof cartridge, removed from the camera, and taken to a photo lab for development. A typical roll of 35mm film contains a length of film sufficient for the exposure of either 24 or 36 images.

Other film formats exist, although they are much less popular than 35mm, especially among ordinary consumers and amateurs. One other format that has enjoyed a moderate amount of success among amateurs is APS (Advanced Photo System), which uses much smaller cartridges than 35mm and is much more automated, making the format well-suited to consumers who prefer not to tinker with their cameras loading and unloading film. Other formats, used mainly by professionals and a handful of dedicated amateurs, include 6x6 (cm), 4x5 (inch), and 8x10 (inch) plate and roll film. They provide superior image quality but are much less convenient, more awkward, and more expensive than 35mm. In addition, most one-hour film laboratories only accept 35mm and APS film for development, so film in other formats must be sent to special laboratories with appropriate equipment.

Film Types

Film can be divided into two broad types: *negative* film and *positive* film. Negative film contains an inverted image after development; that is, the image on the film shows dark areas where the original image was bright, and *vice versa*, and (in the case of color film) shows a similar reversal of colors. Negative film is intended for use in preparing photographic prints on paper, rather than for direct viewing. Positive film yields a directly-viewable image right on the film; it is used for *transparencies* (*"slides"*) intended for direct projection onto a screen or direct viewing in some sort of slide viewer. Positive film is often called *reversal, transparency,* or *slide* film, and negative film is often called *print* film.

Another major type distinction in film is that between *color* film and *black-and-white* (*B&W*) film. Color film produces images in color, just like the original scene being photographed, whereas B&W film produces images in shades of gray, showing only the relative lightness and darkness of details in the image, without any color. Today, color film is far more popular than black-and-white film, although the latter is still preferred by many photographers for purposes of art and journalism. Color film exists in both positive and negative form, but black-and-white film exists only in negative form (there is no such thing as black-and-white slide film).

Development

The images recorded on exposed film are invisible, or *latent*. They are made visible by immersing the exposed film in a sequential series of chemical baths, in a process called *development*. These chemicals cause changes in the film that convert the invisible, latent images to visible images. This conversion is permanent, and once it has been carried out, the film is no longer sensitive to light, and will remain essentially unchanged for decades. Development is a process separate from the actual taking of a picture and is carried out long after the film is removed from the camera. Most photographers take their exposed film to commercial photo labs for development, and do not develop the film themselves (which is messy and complicated, particularly for color photos).

After development, positive (transparency) films can be viewed directly, or they can be projected onto a screen by a slide projector. Negative films look odd when viewed directly, however, and so they are usually used to prepare photographic prints. The negatives are projected onto a large sheet of light-sensitive paper by a device called an *enlarger*, and the paper (which contains an emulsion similar to that on film itself) is developed to convert a latent image produced by the enlarger into a visible, permanent image. The process of printing reverses the image, so that the negative image on the film comes out as positive on the print. Printing, like development, is usually handled by a commercial lab, although some photographers do their own printing, especially if they work in black-and-white (color printing is much more complex).

Today, many film images are converted to digital form after the film is developed. This is accomplished by using a film scanner that slowly scans the entire frame of film and converts the image on the film to digital data that can be processed by a computer. This scanning step is mandatory when images are to be using in print or Web publishing, because both of these domains are now fully digital. Film scanners range from tiny desktop models for the home to hugely expensive desk-sized devices, and the can scan both negatives and transparencies. The smaller ones are often designed for 35mm or APS film only, but larger models can handle almost any type of film.

Occasionally photographic prints are scanned for conversion to digital form, but this is increasingly rare, since the extra step of printing from the negative (rather than just scanning the negative directly) introduces more opportunity for mistakes and errors that can degrade image quality. Print scanning is advisable only when the original negatives or transparencies have been lost.

Film Speeds

All films do not show the same sensitivity to light; some are much more sensitive than others. The degree of sensitivity to light possessed by a given brand of film is called its *speed*, a reference to the fact that the more sensitive the film is, the faster the shutter speeds with which it can be exposed, and thus the “faster” the film is considered to be.

The sensitivity of film to light is designated using an *ISO rating*. The ISO rating is a simple number that indicates the film’s sensitivity. The ISO scale is linear; that is, a film rated at ISO 400 is four times as sensitive to light as a film rated at ISO 100.

A film is usually considered to be of low speed if it has an ISO rating below 100. Such films work well in direct sunlight but rapidly prove to be too slow for practical use in lower levels of light—they are particularly useless indoors and at night, unless a flash unit or heavy artificial light is used to illuminate the subject. The advantage of low-speed films is their very fine grain (see below).

Films of speeds between ISO 100 and ISO 400 are considered to be of medium speed. The films at the lower end of the range are excellent for outdoor daylight use, while still remaining usable in lower lighting conditions. The films at the upper end of the range are good for use on cloudy days and in shadow, and can readily be put to use at night and indoors, with a bit of extra light.

High-speed films have ISO ratings beyond 400, and particularly beyond 800. These films are used mostly for shooting photographs by ambient light (that is, without any added artificial light for the purpose of taking photographs) in poor lighting conditions. Some films of this category have ISO ratings in the thousands.

Grain

Film emulsions contain microscopic, light-sensitive crystals—usually compounds of silver and elements such as bromine (because many compounds of silver are sensitive to light). These crystals vary in size with the brand of film, and particularly with the speed of the film. High-speed films use larger crystals than low-speed films. When films are enlarged for printing or projection after development, it is possible for individual clumps of crystals and their products to become visible in the form of a sort of sandy texture that covers the image in the photograph. This texture is called *grain*.

Grain is usually considered undesirable in photography, except for artistic effects. Unfortunately, increasing a film’s speed (by increasing its sensitivity to light) invariably increases its visible grain. This is one reason why films are manufactured in a variety of ISO ratings: where fine grain is required, extremely sensitive films (with high ISO numbers) won’t do. Some slow films (ISO 25-50, or below) have such

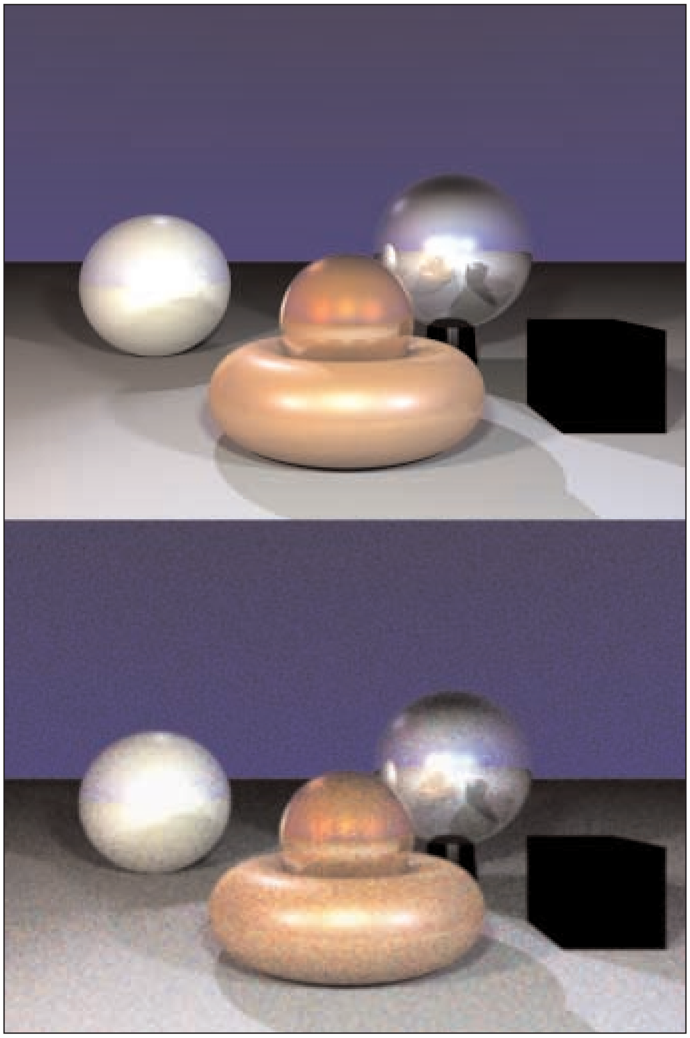

Figure 11

fine grain that it is invisible even after enormous enlargement; but most films of medium speed and beyond will show grain with enough enlargement. It is up to the photographer to decide on the trade-off between grain and film speed that she prefers.

Figure 11 shows an example of film grain. The top photograph contains no visible grain, as would be seen with a low-speed, fine-grain film; the bottom photograph contains obvious grain, as would be seen with a very high-speed film.

Color Balance

Color film emulsions are designed to react to and represent colors in specific ways, based on the type of light that will be used to illuminate the subjects they will be used to photograph. This is called *color balance*. Most color films are balanced for daylight; that is, they are balanced such that they will produce transparencies and prints with appropriate colors if they are used to photograph subjects lit by sunlight (directly or indirectly). Some color films are balanced for incandescent light—the kind of light that comes from an ordinary light bulb. They are said to be *tungsten-balanced* (tungsten is a metal used in the fabrication of light bulbs). There are virtually no other types of color balance in current use.

Outdoor scenes photographed with tungsten-balanced film tend to look very blue, since the film is intended for the more yellowish light of incandescent lamps. Conversely, indoor scenes photographed with daylight-balanced film tend to look very yellow, since the film is intended for the bluish light of daylight. Both types of films give unpredictable results when used with other types of artificial light—very often they produce a bluish or greenish tint in the final photographs.

Color balance *per se* doesn't really exist with black-and-white films. However, all black and white films are sensitive to different colors in different degrees, and this can affect the way they render colored, real-world scenes into shades of gray. The differences between films are slight, but noticeable. Most black-and-white films today are said to be *panchromatic*, meaning that they are sensitive to all colors more or less equally, and produce reasonable black-and-white images under normal lighting conditions.

Latitude and Exposure

The "ideal" exposure of film is that which reproduces a neutral gray color with perfect fidelity. This neutral gray, which is standardized as a gray matte surface that reflects 18% of the light falling upon it, is the basis for all metering and automatic exposure systems.

Of course, in real life, photographing neutral gray surfaces is useless. Real-world images contain detail that involves differences in intensity—light and dark areas, in other words. The dark areas of an image are always less completely exposed than the light areas; that's what makes them dark. So real-world exposure of film involves choosing an *overall* exposure for the film that causes all parts of the image to be close to a neutral gray exposure, even if they don't match it exactly.

At this point, the obvious question is: If film can record different levels of intensity, why is exposure

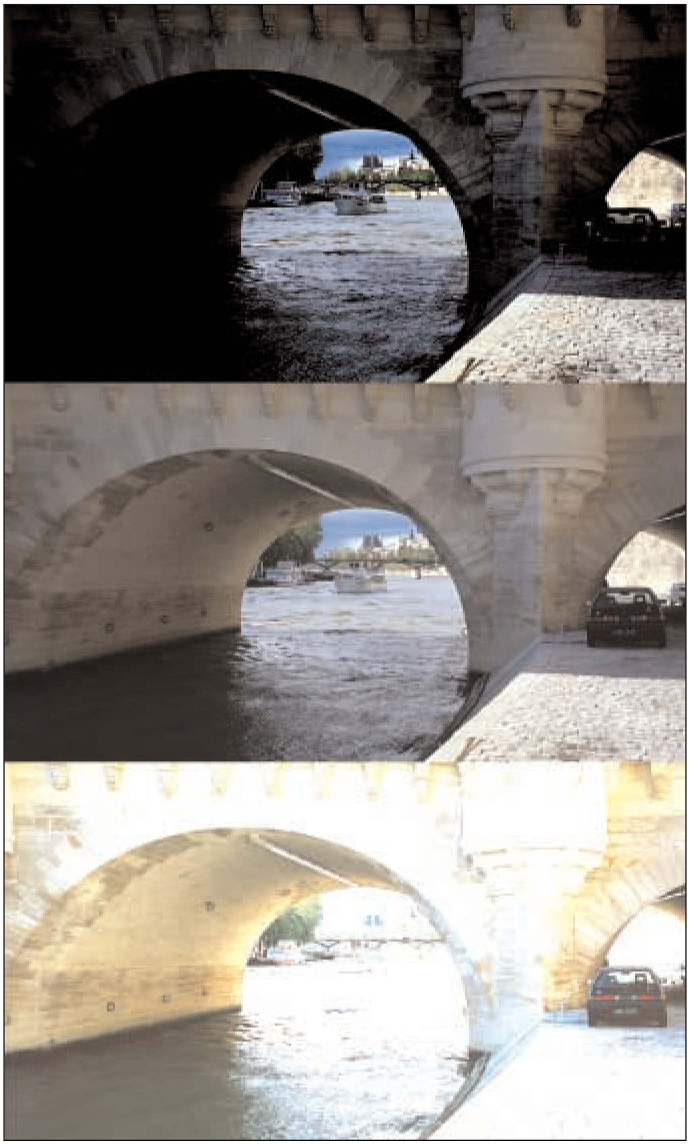

Figure 12

control needed at all? Why not just always expose the film for the same time at the same aperture, and let the dark and light areas of the real image produce the detail? The answer to this is that the range of exposures within which film can be exposed and *still show usable detail* (as opposed to solid black or solid white) is *extremely* limited, compared to the actual range of light intensities encountered in real world scenes. This range is sometimes called the film's *latitude*.

To illustrate the problem further, consider that the brightest parts of the natural world are ten million times brighter than the darkest parts—and yet the average color transparency (slide) film can only handle intensity differences of 32 to one before losing detail! For example, a sunlit day may be 30,000 times brighter than a moonlit night, but photographic film cannot even begin to record an intensity range of 30,000 to one. Because of this, the exposure of film for a sunlit scene must be different than that used for a moonlit scene, in order to ensure that the average exposure of different parts of the sunlit scene is reasonably close to the "ideal" exposure of a neutral gray card. This ensures that as much detail as possible will fall into the narrow exposure latitude provided by the film. On a moonlit night, using the same exposure as that used for the sunlit day will produce an entirely black image, because all of the dimly lit details of the moonlit night fall well outside the latitude of the film for the daytime exposure. Thus, for moonlight, the overall exposure is changed, so that most of the detail in the moonlit scene falls within the film's latitude.

The latitude of a film is thus the difference between the minimum amount of exposure that still allows a bit of detail to be captured and the maximum amount of exposure that also allows detail to be captured. Below the minimum exposure (and thus outside the latitude of the film) shadows turn to solid black, and above the maximum exposure highlights turn to solid white. Figure 12 illustrates this: the center photo is correctly exposed; the top photo is badly underexposed, and most details have faded to black; and the bottom photo is badly overexposed, with most details "washing out" to pure white.

For transparency (slide) films, latitude is fairly narrow: about five stops (a difference in light intensity of about 32 to 1). For negative (print) films, latitude is closer to seven stops (128 to 1). Anything exposed outside this range will come out as either inky black or blinding white, with no detail. These are the values for color films; black-and-white films can have larger latitudes, since they do not need to record color information. No film approaches the actual range of intensities encountered in the natural world around us, however (that would be over 20 stops).

In some cases, the full range of light intensity in a single scene will exceed the latitude of the film, no matter how the exposure is set; that is, the difference between the darkest and lightest parts of the image will be more than the film's latitude of five or seven stops. When this happens, it will be impossible to expose everything in the image correctly, and some detail will inevitably be lost, either in the shadows or in the highlights, depending on the actual exposure value chosen. Many dramatic differences in light level that are easily discernable to the human eye cannot be accurately captured on film, because the latitude of film is far smaller than the adaptability of the human eye. The difficulty in capturing images on film that retain detail in all parts of the image is sufficiently great that many photographers have devoted themselves to developing systems for calculating exposure and manipulating prints and film to compress real-world scenes into the limited latitude of negatives and prints. One of the most popular systems of this type is the Zone System, developed by photographer Ansel Adams and others. These systems are far beyond the scope of this paper, however.

Reciprocity Failure

Reciprocity failure (or *reciprocity* for short) is an unusual property of film emulsions that normally is not a problem for photographers; however, it may become a factor for extremely long or extremely short exposures.

Reciprocity failure causes the sensitivity of film to become non-linear for very short or very long exposures. That is, normally, if you want to expose a film twice as much, you just double the exposure, and if you want to expose it half as much, you cut the exposure in half; but in the case of reciprocity failure, the required exposure value to expose the film by half as much or twice as much is significantly different from the half or twice the exposure value that one would expect. Reciprocity failure occurs only for extremely fast or slow shutter speeds.

For example, astronomers—who may expose a single image for many hours during an entire night—must worry about reciprocity, and appropriately compensate in their exposure calculations. However, for ordinary photographs, with exposures between 1/1000 second and several seconds or so, most films show no reciprocity failure, and so no compensation is necessary. One example of a safe range for film is that of Fuji Provia 100F film (a popular color transparency film), which requires no reciprocity compensation between shutter speeds of 1/4000 second and 128 seconds. Outside of this range, it is necessary to consult a special table for the film that gives the adjustments necessary in exposure calculations. Since virtually no photographers use shutter speeds higher than 1/4000 second or 128 seconds, however, most photographers need never care about reciprocity failure with this film. And the same is true for most films in common use, with most types of photography.

Conclusion

This concludes our description of the technical basis for film photography. Of course, this paper has provided only a very simplified overview—the field of photography is complex, and any comprehensive treatment would fill volumes. Our emphasis here has been on the basic process of taking photographs, as seen from the photographer’s viewpoint; and even within that context, space constraints have limited the amount of detail that we were able to provide (*e.g.*, we necessarily limited discussion to the 35mm film format exclusively). However, this paper provides enough information to satisfy much of the curiosity of the novice, and to perhaps encourage her to learn more from other sources of information, and from experienced professionals and dedicated amateurs already familiar with the science and art of photography.

ABOUT THE AUTHOR

Anthony Atkielski is an amateur photographer and DTP fanatic who enjoys finding excuses to write and publish electronic documents. He lives with his cameras and computer in Paris, France.