Exposure: Key Concepts Explained – Simon Pelling

Introduction

There are a few general articles on exposure on the BirdLife Photography web site, under Resources > Our Articles>Photography Techniques and Processes, for example:

- A four-part series by Bob Young (2012);
- Exposure and ISO: What is the best combination by Glenn Pure (2015); and
- Setting Exposure Revisited by Glenn Pure (2019).

I recommend readers look at these articles. In this article, my objective is to provide a walk-through of the key concepts in exposure; what the terms mean, and how they work together in determining the final picture outcome. I have also included information on light metering, although this topic is also covered in other articles.

Defining exposure

Exposure in digital photography is the process of managing the brightness and colour of the final picture.

There are two parts to this:

- Managing the amount of light (number of photons) hitting the camera's sensor by controlling the **aperture** of the lens, and the **shutter speed** of the camera
- Controlling the 'sensitivity' of the camera, which involves the amount of brightness added through the camera's internal processing (**ISO**).

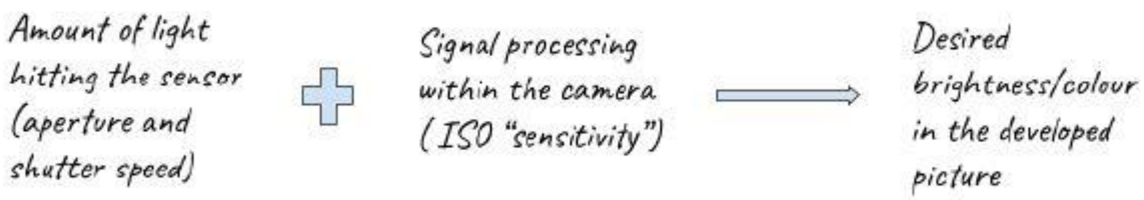

Together, shutter speed, aperture and ISO are sometimes referred to as the 'exposure triangle' even though ISO has nothing to do with the exposure of the sensor to light (see ISO section below).

Why do you need to control exposure?

In simple terms, the pixels on the sensor of a digital camera are stimulated by light (photons) and generate a small electrical signal. This signal is converted to digital data which is used to produce the final picture. However, pixels have a finite ability to respond to light. With too much light pixels become 'full' (or saturated) and adding more light will not increase the signal they produce. With too little light, they will not produce an adequate signal to provide useful information. The purpose of managing exposure is to ensure that enough light hits the sensor to give a picture with the desired spread of tones, from black to white (or dark to light).

Stops and Exposure Values (EV)

A *stop* is a step change in either aperture or shutter speed resulting in (approximately) the *doubling or halving* of the amount of light hitting the sensor.

In addition, the term is used for a change in ISO that has the equivalent effect to a stop adjustment in aperture or shutter speed on the brightness and tonal ranges of the final picture.

Defining stops in this way is useful because it means that a one-stop adjustment to any of the three parameters (shutter speed, aperture and ISO) will have about the same effect on the photograph. For example, an increase in shutter speed of one stop will halve the light hitting the sensor, which is equivalent to reducing the aperture by one stop.

Because it is defined in this way, the stop becomes the 'universal language' of photographic exposure - a stop means the same thing, regardless of the camera and lens being used.

Another, and perhaps more intuitive, term for a stop is '*exposure value*' or EV. One EV is the same as one stop; for example, increasing, say, aperture by one stop will increase the exposure by one EV (which will double the amount of light hitting the sensor).

Cameras usually also allow adjustments of half or one third stops/EVs to allow more fine-grained adjustment of exposure.

Shutter speed

The shutter speed is the amount of time that the shutter is open, in order to allow light from the scene to hit the camera's sensor. Because it relates to a time period, rather than the speed of any moving components, it is perhaps more accurately described as '*exposure time*'. This can be achieved by a mechanical or electronic shutter.

- Mechanical shutters have a blind which opens and closes to expose the sensor to light.
- Electronic-first-curtain shutters combine the electronic activation of the sensor with a mechanical shutter. The exposure is started (with the mechanical shutter in the open position) by reading the data off the sensor (generally row by row of pixels). It is then ended by physically shutting the shutter blind. EFC shutters can be used by mirrorless or compact cameras, or DSLRs in live view, where the mechanical shutter is open by default to allow the viewfinder and/or screen to operate.
- Fully electronic shutters operate entirely by activating and deactivating the process of recording data off the sensor. This option is available in many advanced mirrorless or compact cameras and is the default in smartphones (which have no mechanical shutter).

There is a standard set of shutter speeds used in cameras, which represent one-stop changes. For example, a speed of 1/125 second means that the sensor is illuminated for 1/125 second, or 8 milliseconds. A speed of 1/250 second halves the amount of light hitting the sensor compared to 1/125 second, and a speed of 1/60 second doubles the amount of light compared to 1/125 second.

The standard one-stop set of shutter speed values normally used is (in seconds):

1	½	¼	⅛	1/15	1/30	1/60	1/125	1/250	1/500	1/1000	etc

Modern cameras often have shutter speeds of 1/4000s or even higher, and will allow exposure times of 30 seconds or more.

Cameras typically have available half-step or one third-step values in between this standard set; for example, between 1/500s and 1/1000s you will normally find 1/640s and 1/800s.

For convenience, cameras tend not to display the shutter speed as a fraction, but rather just as a whole number; for example, your camera's display might show 1/125s simply as '125'.

Cameras typically also have a shutter speed represented by 'B'. This is short for 'bulb' which is a setting where you can hold the shutter open for as long as you want, as long as the shutter button is pressed. The name 'bulb' comes from the time (in the early days of film photography) when the photographer would squeeze a pneumatic bulb to hold the shutter open for long exposures.

Photographers use shutter speed to control *the effects of movement* on the photograph. Fast shutter speeds are needed to freeze movement by the subject, or to compensate for the camera shaking. Slow shutter speeds can be used when the subject is static, or to give deliberately blurred effects, such as blurred water in a stream. Bird photographers often prioritise shutter speed in setting exposure because birds are often moving.

Aperture

Lenses have an adjustable diaphragm which creates an approximately circular gap or iris through which light travels. The larger the area of this gap, the more light travels through.

The photograph shows my old manual 50mm lens showing the aperture, formed in this case from six thin metal blades resulting in a roughly hexagonal hole. With most modern lenses the aperture mechanism only closes down the instant the photo is snapped.

The amount of light reaching the sensor depends on the area of the aperture *relative to the focal length of the lens*. The longer the focal length of a lens, the larger the aperture gap needs to be to get a particular exposure.

This accounts for the odd way of describing apertures. Apertures are written as *fractions* of the focal length (f) of the lens. This is called the *f-number*. Readers will be familiar with the f-numbers displayed in their camera - for example, f/5.6 (f divided by 5.6) or f/8 (f divided by 8). Sometimes cameras drop the 'f/' in their displays to save space, simply reporting aperture as 5.6 or 8.

By expressing apertures as a ratio in this way, the f-number is effectively standardised in a way which is independent of the focal length of the lens. A 50mm lens at f/8 will deliver approximately the same amount of light to the sensor as a 400mm lens at f/8, even though the actual diaphragm gap diameters are quite different (6.2 mm for the 50mm lens i.e. 50 divided by 8, and 50 mm for the 400mm lens i.e. 400 divided by 8).

But why use the slightly odd sequence of number values like f/2.8, f/5.6 etc? These comprise the sequence of numbers that are obtained when you calculate the aperture area needed for single stop jumps, starting at f/1.0. In other words, the light reaching the sensor doubles (or halves, in reverse) with each step.

The standard sequence of f-numbers obtained in this way is (showing up to f/22):

f/1.0	f/1.4	f/2.0	f/2.8	f/4.0	f/5.6	f/8.0	f/11.0	f/16.0	f/22.0

For the mathematically minded, this is the sequence derived from $[\sqrt{2}]^n$ with numbers rounded to one decimal place (from left to right, each number in the sequence represents the previous number multiplied by 1.4).

Note that as the numbers are fractions, the larger the dividend (the number to the right of the slash), the smaller the actual aperture. f/8 is a smaller aperture than f/4.

Camera manufacturers usually also use intermediate steps between these standard values, typically one-third of a stop. For example, my Canon camera uses f/6.3 and f/7.1 between the f/5.6 and f/8.0. This allows more fine-grained exposure adjustments than the standard scale.

> Photographers use aperture to control *depth of field* and *sharpness* in images. Small apertures have greater depth of field, which means more of the subject is in focus from front to back. However, if the aperture is too small, diffraction of light introduces blurring and reduces sharpness. Most 'pro' level lenses reach optimal sharpness within a stop or two of the widest aperture. Thus there will be a compromise between depth of field and absolute sharpness.

ISO

ISO is the short-form name the International Organisation for Standardisation has adopted for itself (it's not an acronym, any more than Ted is an acronym for Edward). Standards developed by this organisation are given the prefix 'ISO' and cover a vast range of technical, engineering and other topics. For example, ISO 8601 sets the international conventions for date and time formats, and ISO 13266 sets specifications for attaching ISOFIX child seats in cars. Standards such as this help ensure product safety, quality and interoperability.

Several ISO standards apply to film and digital photography.

With film, these ISO standards relate to the sensitivity of different types of film to light. Films are given a different ISO value depending on the response of their light sensitive coatings to light. The standardisation of ISO ratings meant that a standard approach to exposure could be implemented in mass-produced cameras; a film with a particular ISO rating would always respond to light in (approximately) the same way, so the exposure systems of cameras could be standardised across different systems and brands.

With the development of consumer digital cameras, a similar approach to ISO ratings was implemented. The intention was to create a system which was consistent with film ISO, so that a digital camera operating at (say) ISO 100 could produce a picture that has about the same spread of tones as ISO 100 film when exposed to the same amount of light.

Digital sensors have a fixed sensitivity - they will always respond to photons of light in the same way. ISO does not change the sensitivity of the sensor. Instead, ISO settings are implemented through processing of the sensor signal in the camera *after* the sensor has been exposed to light. With film, if you want to change the ISO you have to change the film, but with digital, you simply turn a dial.

The actual way in which ISO settings are implemented in a digital camera is quite complicated - and indeed is apparently not applied in the same way by all camera manufacturers. However, the key part of ISO settings is *signal amplification*. Essentially, the electrical signals coming off each pixel in response to light can be 'turned up' in the camera - analogous to turning up the volume in a music player.

Thus, changing the ISO is essentially a way of compensating for a lack of light hitting the sensor. If the light reflected from the scene is bright, the signal will require little if any amplification to produce a well-exposed picture. However, if the light is dim, the sensor signal will need to be amplified to produce the required output. Different ISO settings affect the amount of amplification.

The numbers in ISO (eg ISO 100, ISO 200) are a numerical scale of units of sensitivity. This scale was developed for film photography, and carried over into the digital environment. The basic scale adopted for most cameras uses ISO 100 as the base unit, with each step doubling the value; i.e.

ISO 100	ISO 200	ISO 400	ISO 800	ISO 1600	ISO 3200	ISO 6400

Each of these numerical steps constitutes the adding of one stop of brightness to the final image. In other words, changing the ISO by one stop, such from 100 to 200, is equivalent to increasing aperture by one stop, or using a one stop slower shutter speed.

Most cameras also use intermediate values between these basic steps eg ISO 320 or ISO 640 which are approximately equivalent to half-stops.

ISO affects the JPEG output of the camera directly - indeed, ISO settings are primarily about producing a final JPEG image in-camera. ISO settings also are applied to raw files, but the photographer is in the hands of the raw processing software developers as to precisely how ISO (and other) camera settings are implemented when processing raw files. Raw processing software will apply brightness adjustments from information in supplementary data (called *metadata*) included with the raw file by the camera, to produce the basic image for editing. A key difference, of course, is that raw software will apply its own processing algorithms to the raw image file, so the results may not look the same as the JPEG out-of-camera photo.

ISO and noise

Noise in digital photography is essentially slight variations in the brightness and colour information from pixels, giving a digital picture a grainy or blotchy appearance which is particularly noticeable in smooth areas of colour such as sky. This arises principally from the randomness in the way that photons of light strike a pixel. One pixel may receive more or fewer photons during an exposure compared to an adjacent one, resulting in a slightly different output. Other sources of noise include the circuitry of the camera.

Unfortunately, the ISO amplification in the camera does not distinguish between the desired sensor signal and the noise. In low light situations there is a comparatively high level of noise relative to desired signal, so increasing the ISO to high levels creates a bright but very noisy image.

To minimise noise, it is always better to increase exposure by increasing the amount of light hitting the sensor (by increasing exposure value) rather than increasing ISO. Of course, in some situations obtaining a noisy high-ISO image is preferable to getting no image at all, and modest levels of noise are relatively easy to rectify using modern noise reduction software.

Measuring the exposure - understanding the light meter in your camera

- *Middle grey*

Light meters in digital cameras are calibrated to give a reading which is aimed at ensuring that the mid-range tones in an image are properly exposed. This will make maximum use of the sensor's 'dynamic range' - the range of tones from black to white that can be recorded. This mid-range tonal value is known as '*middle grey*'. This has a specific technical definition, but in simple terms it's a shade of grey which is about halfway between black and white in brightness.

When you point the camera at a scene, the light meter measures the brightness of the reflected light and indicates the exposure setting needed to achieve middle grey. The photographer, or the camera (when in automatic mode), will select the desired combination of aperture, shutter speed and ISO to achieve this exposure. This generally works well for scenes with a wide range of tones, such as an evenly-lit landscape, where metering for middle grey would result in a good exposure with optimal range of tones between black and white.

- *Problematic exposure scenarios*

Exposure for middle grey can sometimes give incorrect results.

Scenes with a high percentage of very light tones: The classic case is a snow scene, where the image is heavily dominated by white. The light meter will indicate, as the correct exposure, a value which is aimed at achieving a middle grey tone. Therefore, using this measurement to set your exposure will result in the white snow having a grey tone. The image will be significantly *underexposed,* and you would need to dial in an extra stop or two of exposure to compensate.

Scenes with a high percentage of very dark tones: Likewise, if you shot a predominantly black or dark grey scene, faithfully following your light meter reading to set exposure would result in an overexposed image, with lightening of the dark tones. To preserve the dark tones, you would need to underexpose the photo.

Scenes in which the subject is significantly darker or lighter than the general background: Two typical bird photography examples are a bird against a bright sky, or a white bird (such as an egret or cockatoo) against a predominantly dark background.

In the first example, the exposure value from metering across the scene would be predominantly based on the sky, meaning that (if used) it would result in the bird being significantly underexposed. In this case, the photograph must be overexposed (often by as much as two stops) to get the bird correctly exposed. In the second, the meter would read the scene predominantly as a dark background, and compensate to brighten it, meaning that the photograph would have a very bright, white bird with all detail washed out. To get the bird correctly exposed you would need to underexpose the image.

- *Metering modes*

It is important to understand the different metering modes in the camera. These change the priority given to specific parts of the scene in terms of metering, and will therefore affect the application of the exposure reading to the image.

Usually cameras have at least three different meter settings which change the area of the scene assessed, which are useful in different scenarios. I recommend that readers read their camera manual carefully (or research online) as there are differences between makes and models in how these meter settings work at a detailed level, and more advanced cameras may allow the user to change the operation of the meter. For example, metering may or may not be affected by the position of the focus spot in the frame.

Spot metering: This mode uses a single central spot in the viewfinder, and ignores the rest of the image. This spot must be positioned over the part of the image on which the exposure is to be based. This is particularly useful where there is a specific part of the image to be given priority. However, it can be quite difficult to position the spot exactly, particularly at a distance, and where there is camera or subject movement. Missing the desired spot can result in a completely wrong exposure outcome - simply shifting the spot from a well-lit area to a moderately shadowed area can change the exposure by a couple of stops. It is generally recommended that spot metering be used with manual mode, or with the 'exposure lock' function, so that the desired exposure can be set and fixed (and if necessary checked by a test exposure), before composing the scene.

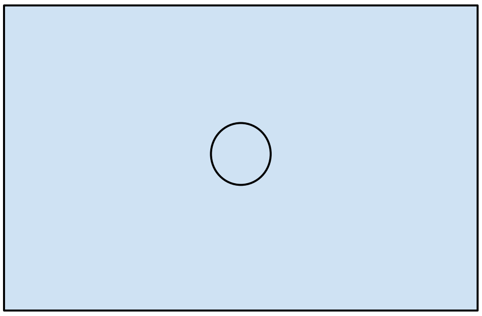

Spot metering

As spot metering only selects a small part of the image, it will often be metering for a single (or very dominant) colour and brightness - rather like the snow scene I described above. As such, the photographer will need to adjust exposure to correct for the fact that the metering will be giving readings for middle grey (eg when spot metering on white birds).

Centre-weighted metering: This metering mode meters across the whole viewfinder area, but gives a greater emphasis on the light towards the centre of the image, on the assumption that (at least in most cases) photographers place the key subject of the photograph at or near the centre of the image. This is a very useful mode for general photography, and it generally gives predictable results. Exposure compensation may still need to be applied where the subject is a different brightness from the background, even if positioned in the centre of the viewfinder.

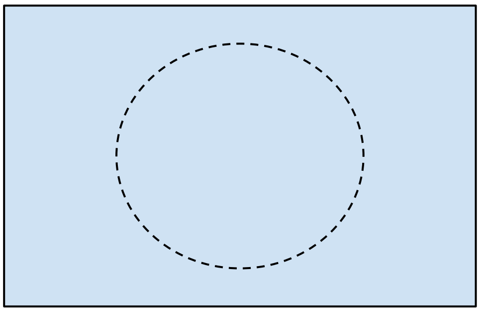

Centre-weighted metering

'Smart' metering: The term 'evaluative metering' is used by Canon, and other manufacturers will have something similar but under a different name, such as 'matrix metering'. These smart systems probably differ slightly from camera brand to camera brand. These modes divide the area of the viewfinder into a number of zones (for example, a matrix of 7x9=63 zones). Metering is still undertaken across the whole frame, but the camera measures the light in each zone separately, and uses some means of determining the priority to be given to each zone in determining the overall exposure.

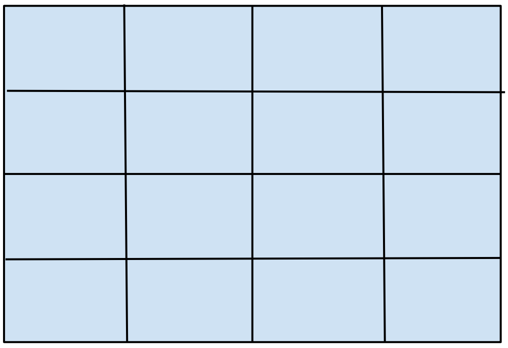

Evaluative metering

The precise means used by different camera systems to assess zones are complex and often not disclosed. I understand that they commonly use focus points; that is the camera gives extra weighting to zones that include focus points that are in focus, because it is assumed that the photographer's highest-priority subject or subjects are those that are focused on. However, cameras also use other information such as colour and pattern matching (using a stored database of different images) in prioritising zones.

Evaluative modes are good for general photography and usually work well, and are often the default metering set in the camera 'out of the box'. However in my experience these modes are often no better than centre-weighted metering in many bird photography scenarios - that is, the 'smarts' don't necessarily give a better outcome in bird photography.

- *Reading the meter*

Setting the correct exposure involves the use of a marked scale in the camera viewfinder (or rear screen) that shows when the aperture, shutter speed and ISO are together delivering the metered reading. The scale is usually marked in one-stop steps, with half or one third steps in between. A moveable point is matched against the scale to show the effect of exposure adjustments.

When the moveable point is positioned in the centre of the scale the camera will be set to expose according to the metered reading. In the problem exposure scenarios set out above, changes to exposure will then need to be 'dialled in' moving the moveable point to the right (for over exposure) or to the left (for under exposure) of the centre of the scale.

In manual mode this is straightforward; the photographer just adjusts the desired parameter (shutter speed, aperture or ISO) to position the moveable point at the desired position on the exposure scale. However, if you shoot in a more automatic mode (i.e. settings where the camera chooses at least one exposure parameter, such as Aperture Priority), the camera's computer will base the settings it selects on the metered light measurements. In problematic lighting, you would therefore need to force it to shift its settings.

This is known as *exposure compensation*. Digital cameras normally have a button or dial settings which allow you to dial in positive or negative Exposure Values (EV) as compensation. For example, to force the camera to overexpose by one stop, you would press the relevant button and then dial in +1 EV on the scale in the viewfinder. The camera will change one or more of the three exposure parameters (shutter speed, aperture and ISO) to give effect to the change.

Mirrorless cameras use electronic viewfinders and these are normally set so that changes to exposure values are reflected in the view on the viewfinder screen; adjust the exposure a stop brighter, and the view through the viewfinder will also become brighter by an approximately equivalent amount. With experience, advanced photographers often learn to trust the viewfinder view to set the desired exposure, rather than relying on the scale being shown and then estimating the amount of exposure compensation required for difficult lighting. Exposure settings are changed until the subject appears properly exposed in the viewfinder. Often a histogram can be displayed in the viewfinder for added confidence.

DSLRs have a direct optical view of the scene so this approach is not possible when using the viewfinder. However, a similar experience can be had with the rear screen in DSLRs in live-view, although using the rear screen is often difficult with long heavy telephoto lenses when hand-holding the camera.