

THE CAMERA IS AN IDIOT – WHY YOUR CAMERA LIES TO YOU AND WHAT YOU CAN DO ABOUT IT

Our cameras are recording devices, capable of capturing an incredible amount of data, but they are very poor at communicating that to the operator. The image shown on the camera back is inaccurate, the exposure reporting is unreliable and the colours at best patchy. Why is this so?

INTRODUCTION

When taking photographs in the field (or more correctly “**capturing data**”) we rely on the information that the camera displays, including the thumbnail on the camera back, the histogram and the exposure readings. **The problem is this information is not necessarily accurate and may in some cases be downright misleading. There are at least 5 things that our cameras do that drive us crazy. These are:**

1. The image on the back of the camera is not the RAW file and can look very different from what we captured, making any sort of evaluation difficult.
2. The camera often indicates highlight clipping, except there isn't when we open the file in Software.
3. We often see distracting elements in the background when we open the image, but these weren't noticeable when we set up the shot.
4. The light meter doesn't always make the right choice of exposure, but because of problems 1 & 2, we have trouble assessing if it is OK.
5. Sometimes the White Balance (Colour Temperature) is way off when the image is reviewed and opened in software.

Items 1 & 2 are related, so let's look at them together, before we look at the other issues:

1. THE IMAGE ON THE CAMERA BACK LOOKS DIFFERENT FROM THE RAW FILE AND

2. CLIPPING WARNINGS ARE INCORRECTLY REPORTED BY THE CAMERA ON THE DISPLAY

The problems with the image we review on the camera back stems from a series of issues relating to how the camera displays the captured information, which can be inaccurate in several critical areas, including:

1. The image review on the sensor is **NOT the RAW file** that we capture, but a lower resolution JPEG processed in camera using the manufacturer's settings. Your RAW file may be 24, 36 or more Mega Pixels (Million Pixels), but the screen on your camera has only about 2 million dots, so it is not possible to display the full size file. In addition, there are many different ways to process a RAW file, so it is misleading to present one version, that may or may not be how you want to interpret the file.

2. The histogram is based on the JPEG render and therefore does not represent the full data captured. A JPEG file is 8 Bit only and can only be in either sRGB or Adobe RGB colour space, your RAW file is 12 or 14 bit and has no colour space attached. Part of the JPEG conversion includes data compression and application of a tone curve, that will clip some highlight and shadow detail.
3. The colours shown on the screen are optimised for the JPEG and not the RAW file, probably mapped as sRGB to suit the monitor and to make the image "look good". They do not represent the "real" colours in the processed RAW file (in actual colour represented, or the colour gamut).
4. The white balance used to render the JPEG will be what the camera recorded, the RAW file does not have any white balance set (when opening in software with default settings it will display the White balance set for capture, but this is not fixed), so at best, the white balance is a guess.
5. The resolution and colour accuracy of the screen on the camera back is not as good as a studio monitor and can't display the full colour gamut correctly, so even if the full colour detail was available to display, you wouldn't be able to see it anyway.
6. The exposure clipping warnings are also based on the JPEG compressed file and are likely to show clipping of highlights when none exists.

So, how much can we rely on the information the camera is giving us and how can we be confident that we are capturing good quality data, ready for processing into an image in our chosen software?

HOW THE CAMERA DISPLAYS CAPTURED FILES

To understand why the camera gives us misleading information, we need to understand what is being recorded and how that is viewed.

When we capture an "image" in RAW format, we are not actually recording a viewable image. The information captured is not viewable without processing applied to convert the RAW data to something we can see.

To provide the user with something that can be viewed on the camera screen, the RAW data is processed to display a JPEG image, usually at lower resolution than captured. This is processed by the camera using settings designated by the camera manufacturer. Generally, we don't have much input into this process, other than selecting one of several Picture Controls (Nikon), or equivalents that allow some adjustment of tone, saturation and contrast.

The RAW data is converted into a JPEG thumbnail using these picture control settings, although due to the JPEG conversion process, the information displayed cannot accurately reflect the RAW data captured.

The reasoning behind this stems from the early days of consumer digital capture, starting around 25 years ago. The only file format available to consumers then was a JPEG or TIFF and each camera manufacturer came up with their own "recipe" for how the image data was converted. In 2004 Adobe released the first version of its DNG format and Photoshop was first able to read and process RAW files in the 2003 release of Photoshop. At that time the range of controls available to process the data were extremely limited and RAW processing was in its infancy.

For comparison, the first "serious" Nikon DSLR was released in 1999, but the earliest consumer digital cameras were available from the early 1990s. These only supported JPEG or TIFF capture. It was not until software was developed that could process RAW data that RAW capture started to become available on digital cameras. Therefore, the cameras were designed to display a JPEG on the rear screen.



Above: the first Nikon DSLR the D1, 2.7 MP

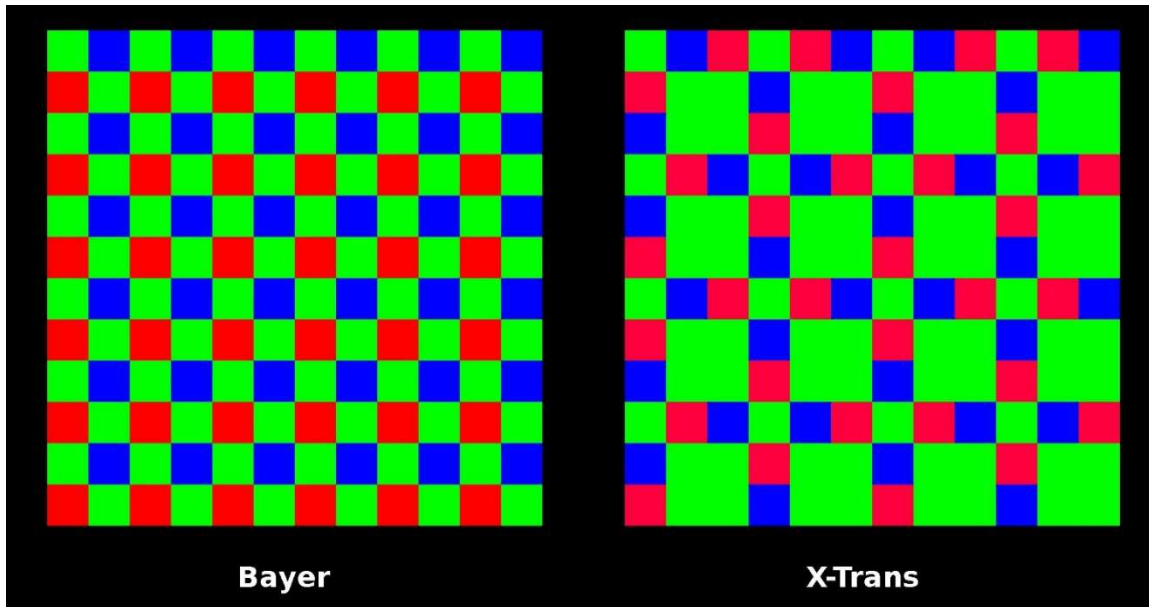


Above: An early Nikon Coolpix, 3.2 MP

Because processing and displaying a RAW capture requires considerable processing power, and there are a multitude of ways the data can be interpreted, camera manufacturers have stuck with the "tried and true" method of displaying the image based on a JPEG conversion.

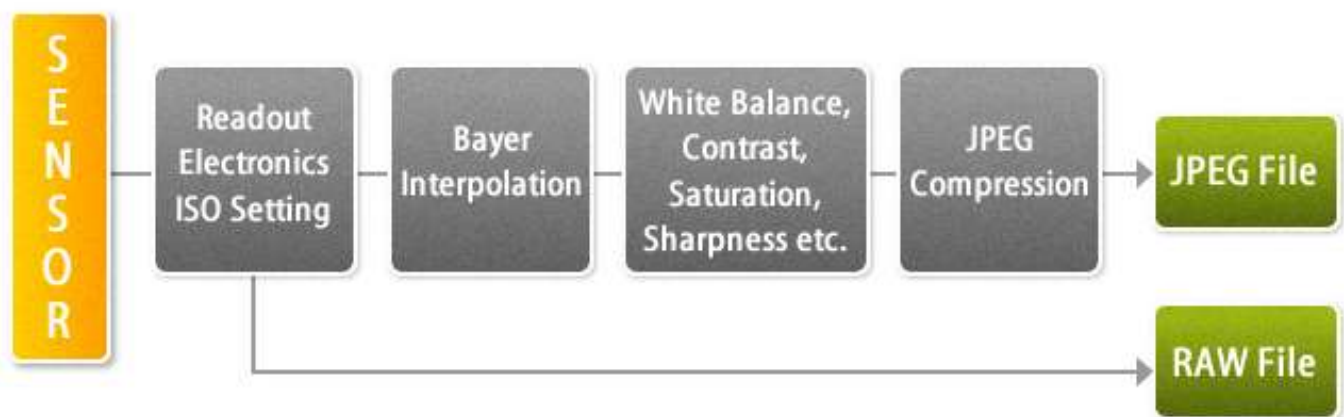
HOW CAMERAS CAPTURE DATA

Cameras capture the light falling onto the sensor and the light is filtered through a mosaic of Red, Green and Blue filters covering the individual photo sites (pixels). Each photo site only captures one colour of light, so in its most basic form, the light falling onto the sensor does not make up a coloured image. Each photo site for each colour channel measures light intensity in up to 16,384 possible levels of brightness in a 14 bit capture (or 4096 levels for a 12 bit capture). At this point no colour exists for these readings as the actual colour value will be allocated during the conversion process.



Above: Two typical sensor configurations. Each photo site (pixel) captures only one channel of colour.

To convert the captured data into an image, there are two different paths, one for a JPEG render and the other for RAW processing. The RAW processing can only be done outside the camera.



Above: How the RAW data is converted inside the camera to a JPEG. It should be noted that a JPEG file can only exist in either sRGB or Adobe RGB colour space, whereas the RAW file can be processed in a much larger colour space such as ProPhoto RGB (as in Lightroom) and that is only determined during processing and output as an image file (TIFF, PSD, JPEG etc).

The look of the JPEG image on the back of the camera therefore bears little resemblance to the unprocessed RAW file and can only be regarded as a very rough guide. It is primarily designed to look "good" on the camera back. In addition to the "sins" committed in turning the RAW data into a JPEG, the bit depth is also compressed from either 12 or 14 bit to 8 bits, meaning a maximum of 256 levels of tone. To achieve this a lot of data is lost and the highlights and shadows compressed.

INACCURATE CLIPPING WARNINGS

Because the image data has been compressed and reduced to 8 bits, some tonal values are lost from the RAW capture. The histogram shown on the camera back is derived from the JPEG thumbnail, so it reflects this compressed information.

That is why the clipping warning on the back of the camera does not reflect reality.



Above: Clipping warnings on the camera back compared to the RAW file opened in Lightroom. The camera shows clipping, the RAW file does not (simulated example).

Often, an image thought to be clipped when viewed on the camera back will display no clipping when opened in Lightroom or other RAW processing software. Lightroom reads the full data from the RAW file.

In summary, it is impossible to accurately “see” the RAW data captured by the camera on the camera back, but we can use experience and a little testing to evaluate how to interpret what the camera is telling us. But first, let’s look at how RAW files are processed in software.

THE DIGITAL “NEGATIVE”


RAW files are often referred to as a “**Digital Negative**”. This is not an accurate representation of the status of a RAW file, it is more like an exposed roll of film that is ready to be “processed” into a viewable image – which then becomes the digital equivalent of a “Negative”. In the same way that a roll of film can be processed in differing chemistry, processing a RAW file can produce a very different image, so it is inaccurate to call this file a “Negative” as it implies some fixed status and ability to be viewed. It does not have either.

RAW DATA PROCESSING

The processing of RAW data into a viewable image includes several steps:

Demosiacing of the R, G and B pixel readings into a grid of pixels that contain R, G & B values for each pixel, based on the ISO setting and the exposure set by the user.

Application of a tone curve to replicate the Gamma of human vision. Essentially as we increase the brightness of an object our eyes do not respond in the same way. Our registration of brightness follows a curve and is not linear. Human eyes are more sensitive to changes in brightness at lower levels and less responsive at higher levels. RAW files are mapped to more closely reflect human vision.



Above: Comparison of human eye response to camera linear recording of luminance.

Rendering of colours in a default editing colour space – for Lightroom it is a version of ProPhoto RGB, which offers the largest usable range of colours. Colour space is not applied until the image is exported into Photoshop, or as any of the file types supported (TIFF and PSD can be any of the colour spaces, JPEGs can only be sRGB or Adobe RGB).

Application of base level noise reduction and sharpening.

Application of default level contrast (as part of the tone Curve)

Application of the “As Shot” white balance – this is not fixed and can be modified during processing.

The default view of a RAW file when first opened in software is most likely based on the camera rendered JPEG file, however as we move into the Editing environment and load the full version of the file, the full data set is read and processed, ready for your input.

GETTING THE EXPOSURE “OPTIMISED”

It is difficult to say what a “correct” exposure is when discussing digital images. We can only really talk about the optimal capture of data that gives us the maximum information available when rendering into a viewable image. The old axiom of “ETTR” or Expose To The Right is a good guide for capturing the most data, however shadow details tend to be overly light.

A RAW file is usually captured as a 12 or 14 bit data set, so exposing to the right and recording as much data as possible leaves the most flexibility for processing.

In addition to Exposing to the Right, in order to accurately render shadows and shadow detail, we should also process the RAW file to reflect the actual shadow appearance – that is, an image exposed to the right might render the darkest areas as grey, so the tonal range needs to be adjusted to rebalance the distribution of tones within the image. The axiom should be “**Expose to the Right, Process to the left**” (Dr Les Walkling).

COMPARISON OF DATA CAPTURE

Exposing to the right maximises the amount of data captured in the highlight area whilst maintaining as much information in the shadows as possible.

Note that most cameras will be capturing an exposure range of 10-12 EVs (or “stops”), depending on lighting, so the minimum amount of data captured in a highly contrasty scene (12 EV steps) would be 8 levels at 14 bit, 2 bits at 12 bit and nil at 8 bits. Therefore, an image captured at 12 or 14 bit and compressed to 8 bit is likely to have little or no detail in the darkest shadows and possibly clipping in the highlights as well, depending on how the image data is compressed during JPEG conversion.

SUBJECT EV RANGE	14 BIT LEVELS	12 BIT LEVELS	8 BIT LEVELS
1 Brightest level	16384	4096	256
2	8192	2048	128
3	4096	1024	64
4	2048	512	32
5	1024	256	16
6	512	128	8
7	256	64	4
8	128	32	2
9	64	16	
10	32	8	
11	16	4	
12	8	2	
13	4		
14 Deepest shadow	2		

Exposing an image without optimising the exposure may result in loss of data in the shadow areas, where there is very little detail recorded. Underexposing and pushing the tones to the right to brighten the image will amplify the lack of shadow detail and enhance the shadow noise, so is not desirable.

In a situation where there is a greater difference in tonal values than around 12 stops, it will be necessary to make several exposures at different settings and then combine them into an HDR render.

HOW TO INTERPRET THE IMAGE ON THE SCREEN

EVALUATING EXPOSURE

Pushing the exposure to the right, given that the display of the histogram is not accurate can cause headaches, as you don't really know when the data becomes clipped. One other issue is that the blue channel in the data recording tends to clip earlier than the other channels and is more prone to noise than the other two channels. The best solution is to set the camera to display the separate R, G and B histograms, this provides more information than the combined luminance histogram.




Above: The camera display with the separate R, G and B histograms.

In order to evaluate the accuracy of the histogram displayed on the camera back, make several different exposures of a subject with a range of tones, including white. Gradually increase the exposure levels whilst recording the appearance of the clipping warnings. Take note of the exposure where this starts to show and then evaluate the same images in your RAW editing software to see where the actual clipping occurs.

By comparing where the clipping starts to show on the camera back and where it actually occurs when processing the RAW file, you can evaluate how much clipping in the display can be tolerated before you lose image data. This is likely to be only a small amount such as 1/3 to 2/3 of a stop but taking these test shots will give you an accurate feel for the metering in your camera. Gaining experience with looking at the clipping warning and seeing how much of the image area is clipped will assist in making a decision in the field whether you need to make another capture with a lower or higher exposure setting.

Some minor highlight clipping in the RAW file when opened in software can be easily recovered in processing, so all is not necessarily lost.



Above: the RAW file opened in Lightroom; some highlight clipping is showing at default settings.




Above: the same image with exposure, Whites, Highlights and shadow adjustments, no clipping.

Note: it is inaccurate to talk about “**Post Processing**” unless we are talking about “**Post Capture Processing**” which is essentially an oxymoron. You can’t process before capture, so we are just talking about “**Processing**” (leaving aside the fact that in some cameras there is a bit of “under the hood” application of some settings, like noise reduction and lens corrections). Essentially, we are talking about capturing data that requires processing in software to render it into an image.

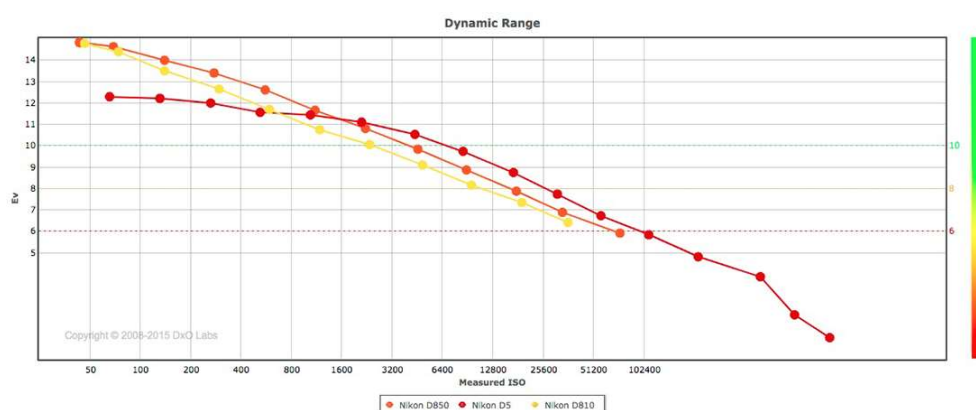
BRACKETING TO THE RESCUE?

Bracketing exposures can be a way to ensure a critical image is rendered correctly, but this is really only useful when using a tripod as the images need to be aligned in the editing program in order to create an HDR render. Many people always bracket their exposures, but in most cases, it is only where there are extreme differences in brightness range and contrast that more than two exposures are needed.



Above: Two exposures bracketed at 1 1/3 stop apart and blended into an HDR DNG (right) One exposure was made to cater for the bright sun, the second was to expose for the shadows. The resulting DNG was adjusted to give the overall tonality required, with no clipping of shadows or highlights.

Another consideration when evaluating exposure is the remember that as you increase the ISO setting on the camera, the dynamic range decreases. Shooting at the base ISO gives the greatest dynamic range, increasing the ISO setting gradually reduces the dynamic range capable of being captured.




Above: Nikon D850 Dynamic range. At base ISO the range is over 14 stops, at 6400 ISO it is 9.5 stops.

EVALUATING COLOUR

There is no accurate way to evaluate the colour in the image on the camera back, as this is based on some assumptions made by the camera JPEG engine, including the camera White Balance settings and the display colour space and compressed colour gamut inherent in the JPEG conversion.

Another factor to consider is that colours in sRGB colour space are not the same colours in other spaces, therefore the overall look of the colour will be different in the editing environment.



Above: maximum Green in sRGB colour space is a long way off Maximum green in Adobe RGB and even further away in ProPhoto RGB (not shown).

In conclusion, the best way to evaluate the colour in the image is to pretty much ignore it and know that the colours in the processed image can be controlled to give the appearance you want in the image.

Let's now move on to the other issues we identified with the way the camera displays information.

3. THERE ARE DISTRACTING ELEMENTS IN THE IMAGE WHEN I OPEN IT THAT I DIDN'T SEE IN THE VIEWFINDER

A common problem when photographing in the field is that when composing an image in the viewfinder, we really can't see clearly things that are not part of the areas we are concentrating on. Part of this is due to the fact that human vision is only sharp in the central area of vision, everything else is blurry. This is compounded by looking through the lens which is wide open and gives us the shallowest depth of field. Using a macro lens, or a longer lens at full zoom, the focus falls off quickly and things in the background are shown well out of focus and so we tend to ignore them.

Depth of field is difficult to evaluate with a DSLR prior to capture as to view the Depth of Field of a scene through the viewfinder requires stopping the lens down to the selected aperture. With small apertures and in low light it can be very difficult to make a proper evaluation as the view becomes very dark.

On a mirrorless camera, we are still looking at the image on the screen that has been converted to a JPEG and at the wide open aperture setting used for composing and framing. Stopping the lens down does not seem to help much as the preview option only gives a couple of stops of preview, so you really can't evaluate Depth of Field with any certainty.

EVALUATING DEPTH OF FIELD AND DISTRACTIONS

The best way to evaluate distracting elements in the background is to make a test exposure and carefully move around the image when zoomed in on the camera back. The thumbnail view gives a good representation of the focus and depth of field, so this is one area where the camera is giving you reliable feedback. If not in a position to make test shots, you have to make an educated guess and try to shoot with the widest aperture possible (lowest number) to avoid background objects “appearing” in the shot.



Above: The viewfinder view (left) is wide open, the actual shot (right) is at the selected aperture, so elements out of focus can suddenly “appear”.

EVALUATING FOCUS, MOVEMENT AND FRAMING

The good part of the thumbnail image that is not “corrupted” by the JPEG conversion is that these aspects of image capture can be fairly well evaluated from the screen on the camera back. Given that the image may be a reduced size version, you should still be able to make a reasonable assessment of these aspects, leaving final evaluation until the full RAW file is opened in your editing program.

Checking the focus point is reasonably accurate, given the probable down sampling of the data, but should be good enough to evaluate. Framing is accurate as is any sense of movement – camera or subject, intentional or otherwise.

4. THE CAMERA'S EXPOSURE METERING IS SOMETIMES NOT ACCURATE

Unrelated to the issues with viewing the image on the camera back, sometimes the camera exposure is just not optimal. Despite advances in technology, special metering modes etc, there are times when the camera metering is just not right.

Essentially, your camera is just a dumb recording mechanism, despite being able to “See” eyes and focus track, the camera really can't recognise anything to do with subject and lighting, other than what it is programmed to do.

Camera metering is designed to give a good result in a range of “normal” situations, but unusual conditions can throw the metering off.

Most scenes tend to have a range of tones and colours, with an average of the lightest and darkest tones being close to mid grey. Most scenes are metered well, but scenes with a dominance of lighter tones will tend to result in an underexposed result as the camera “assumes” the middle tone is 50% grey.



Above: Unusually bright scenes without darker tones will tend to be **UNDEREXPOSED** as the camera's meter attempts to render the middle tone at 50% grey. The metered exposure reading will need to be increased by up to +1 or 1 2/3 stops to render the scene brightness accurately. Take a test shot and evaluate the histogram before shooting again.

Conversely, a scene with predominant darker tones will tend to be **OVEREXPOSED** as the middle tone is adjusted to mid grey. In these situations, reduce the exposure to ensure no highlight clipping.

Scenes with predominately light or dark backgrounds will not be rendered correctly and the central elements can be over or under exposed. Experience with varying conditions should warn the photographer that some adjustment may be necessary.



Above: the bright background has rendered the foreground too dark. Adjusting the settings to increase the exposure without clipping and processing accordingly can give a pleasing result.



Above: the dark background has caused the bird to be overexposed, reducing the exposure renders the tones closer to optimal and avoids clipped highlights. (This is a simulated example)

EVALUATING EXPOSURE

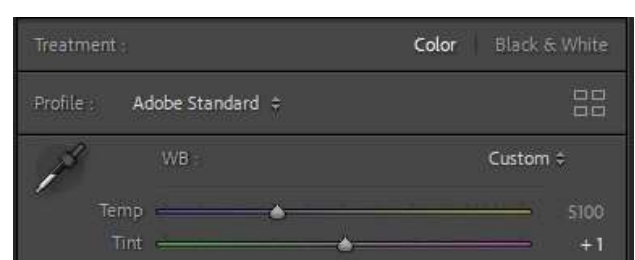

Given that the image on the display screen is not accurate and clipping is not reported well, you might think that you can't make a reasonable assessment of the exposure. However, as we have seen, it is possible to learn how your camera reports exposure and with experience of unusual lighting situations you can make a fairly accurate assessment of how and when to intervene with the camera's proposed exposure.



Above: the camera wanted to underexpose this scene, at right, corrected by increasing exposure. A quick view of the histogram of a test shot showed no highlights at all, so exposure was increased by 1 stop (simulated example).

5. SOMETIMES THE AUTO WHITE BALANCE IS WAY OFF

By capturing RAW data, the white balance is not relevant, however whatever the camera setting for white balance will determine the look of the JPEG thumbnail as well as the default view when opened in software. Setting the camera to "Auto" white balance can sometimes give good results, but sometimes it is wildly wrong. Using any of the presets such as Daylight, Cloudy etc can also give varying results, so I suggest using "Auto" as a consistent starting point.





Above: A forest scene photographed using Auto white balance. The large area of bright green moss has fooled the camera into adding Magenta, at left. Correcting in Lightroom by reducing the Magenta amount restores the "correct" appearance.

ASSESSING WHITE BALANCE

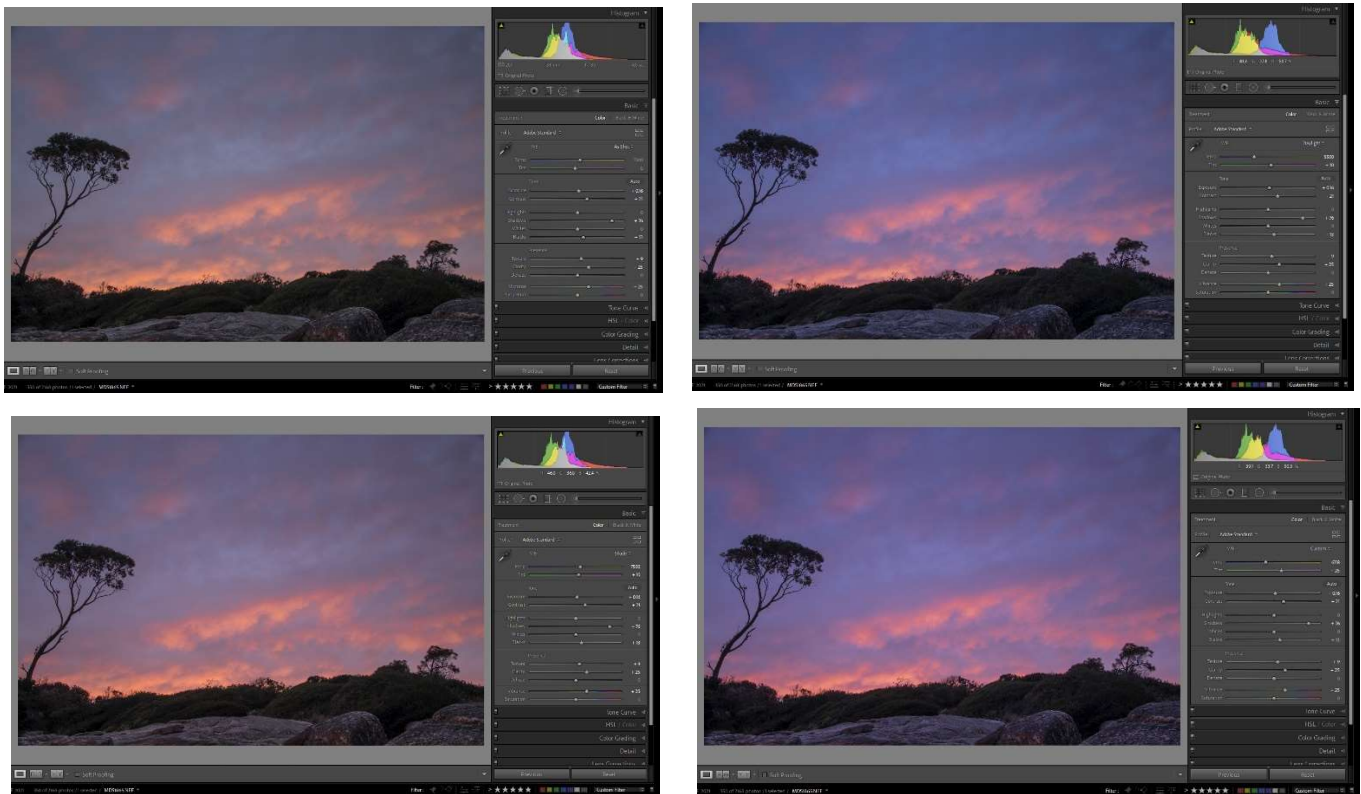
Photographing in the field there are many variables that affect the white balance, such as the time of day, the quality of the light falling on the subject, and the impact of surrounding objects.

Setting the white balance can only really be done in the processing stage using a calibrated monitor. According to Dr Les Walkling, the white balance sets the “mood” of the image and it is entirely a personal preference. The “correct” white balance is what looks right to you, the photographer.

Making any assessment of the white balance of an image using the display on the camera back is therefore a futile exercise, the best option is to just ignore what the camera thinks and only make an assessment from the comfort of your chair in front of the computer.



Above: This kelp was photographed in part shade where the open sky was lighting part of the surface, giving a blue cast. Correction in processing removed the blue and revealed the intended colour.



Above: A sunrise shot with (from top left to bottom right) Auto WB, Daylight, Shade and a custom setting. Which one is “correct”? and does it matter?.

TAKING CONTROL OF THE CAMERA

So, how do you take control of the “idiot” in your hand? As we have seen, the camera does not always tell the truth about what you are shooting, but this has come about through the restrictions on what can be processed in camera and shows how complex things have become with RAW capture and software processing.

So, don’t blame the camera, it is only doing what it has been designed to do. We, the photographers need to use our experience, expertise and creativity to understand the shortcomings of our devices and learn to take charge, when and if necessary.

To summarise what we have found:

Don’t rely on the image on the camera back, it isn’t “the image” just a quick “guess”

Don’t believe it when it says highlights are clipped, they probably aren’t

Don’t always rely on the meter in unusual or extreme cases

Don’t even think about the white balance the camera is showing you, it can be wildly off.

Don’t blame the camera for annoying distractions that you didn’t see when composing.

So, to make the camera smarter, here is my wish list:

1. The screen on the back should show us the RAW file with default processing, plus the option of a choice of colour modes or recipes for us to evaluate in the field. The histogram would then be a more reliable tool. This would require a lot of processing power and take time to load and display. Also, the colours on the screen are not necessarily accurate (and what about the ambient lighting), so the screen would need to be improved.
2. We need a reliable way to preview the depth of field that doesn't involve the screen going dark, and some way of simulating this with mirrorless cameras. This probably isn't achievable on a DSLR as we are looking through the lens. It might be possible for mirrorless cameras with some fancy software "tricks".
3. More reliable White balance, using AI or some other methods, to make the capture settings more usable and predictable. This should be achievable with more "brains" inside the camera.
4. More reliable metering of exposure, to eliminate those "failures". Cameras have gotten much better in recent years, but still can be caught out in extreme situations. More refinement is needed.

INTERNET ARTICLES FOR FURTHER READING

For further reading on this subject, here are two sources:

<https://photo.stackexchange.com/questions/82836/can-you-see-a-raw-preview-with-digital-cameras>

<https://www.cambridgeincolour.com/tutorials/gamma-correction.htm>

© Michael Smyth/Phototutor 2021