

Seven 'Deadly Sins' Found In Our Galleries: Part 2 - Ian Wilson, PhD(optics)

[Editor's Note: This article is a reproduction of Ian Wilson's presentation at the 2015 Digital Photography in the Bush event, at Hunter wetlands Centre, Newcastle. It is being presented in two parts in this newsletter; Part 1 appeared in the June 2015 issue.]

Post-Processing Issues

4. Colour adjustments

Another problem that I sometimes see is incorrect white balance (WB) adjustment which causes images to have an unnatural 'cool' or 'warm' colour cast. It is usually easy to fix so there is no excuse for showing an image with this fault. Most bird photographers shoot with the WB set in the camera on Auto. This setting usually does a pretty good job and there is often no need to do anything in post-processing except to check that the WB is acceptable. To accurately check the WB, one needs to have a neutral white or neutral grey reference standard somewhere in the image; a highly improbable occurrence except in a test scene. However, many scenes have a naturally occurring white reference such as white clouds, breaking waves, snow and the white parts of many birds. You can move your cursor over these white patches and read out the RGB brightness values. If the WB is correctly adjusted, the RGB values will be roughly the same for each colour channel and ideally the brightness of $R = G = B$. If they are not similar in value, and you are satisfied that the white patch is a reasonably accurate white reference, you should use the WB adjustment tool. This is usually an eye-dropper that, when clicked on the reference, corrects the colour cast over the entire image and adjusts the RGB data numbers of the white patch to have equal value. If there is no neutral white patch in the scene, then you may be able to find a patch of neutral grey. Some grey birds, like the Grey Fantail, have plumage that is a good approximation to neutral grey.

WB is best corrected while making the global lighting adjustments described above. If you are processing a RAW file, the task is usually quite easy. If you are unable to find a reliable neutral white or grey patch in the image, it is worth trying the WB appropriate for the lighting in the original scene. For example, if the image was recorded on a sunny day, click on the WB check box for 'daylight' and see if the image looks more like you remember or, if the shot was captured in a shady situation, click on the WB box for 'shade' and see how the image looks. For fine adjustments the 'colour temperature' slider can be used. If you shoot JPEGs and process using Photoshop, you can correct WB problems using the Remove Colour Cast dialogue box. This works well if you have a neutral white or grey patch in the image to use as a reference. However, best practice is to shoot RAW files and avoid shooting JPEGs if at all possible. In Figs 5(a–b) I have deliberately set the WB too cool (daylight) and too warm (shady). It was an overcast day and the correct WB is the in between 'cloudy' setting.



Fig. 5(a). The white balance has been set to 'daylight' making the image appear too 'cool'.



Fig. 5(b). The white balance has been set to 'shade' making the image appear too 'warm'. The correct white balance is 'cloudy' corresponding to the conditions on the day the image was captured.

Saturation and the seductively-named vibrance are related and can be used to increase the intensity of colours. Their main use is to attract attention and for this reason they are much used in the preparation of artwork for commercial advertising. They both need to be used with great care and not in combination. If you have properly adjusted the lighting, there should be little need to use these tools in bird photography. If you are someone who routinely makes global adjustments to the contrast, saturation or vibrance, it is probably time to reconsider what you are doing and why; these adjustments are usually not necessary or beneficial. One of the adverse consequences of making ill-considered global colour adjustments can be surreal blue sky and unreal blue sea. There is a fair range of hue that is acceptable for sky and sea blue but there is a point beyond the pale as for example in fig. 6. In this case I have pulled back the overall brightness one stop to recover detail in the whites and used a global saturation adjustment to increase the strength of red in the bird's bill and feet. An unintended consequence is the unnatural deep blue sky. I should have adjusted the bird and sky separately and used a selective area adjustment on the bill and feet. There is a 'blue sky' standard for digital image reproduction with R= 135, G = 206 and B = 235 DN which we can use as a rough guide. It is not a good idea to go darker than the standard if you want a natural looking blue sky.

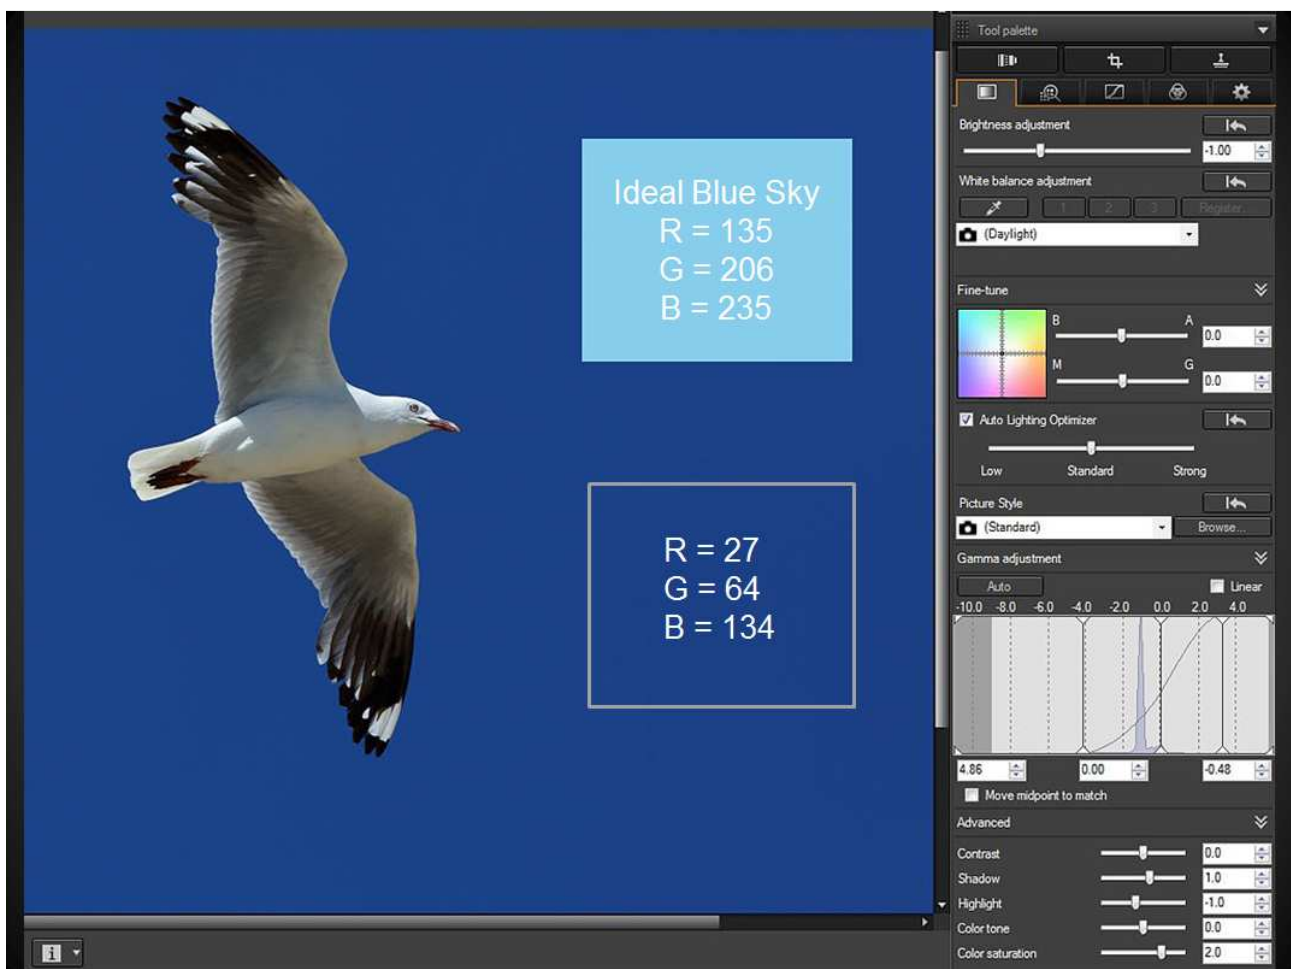
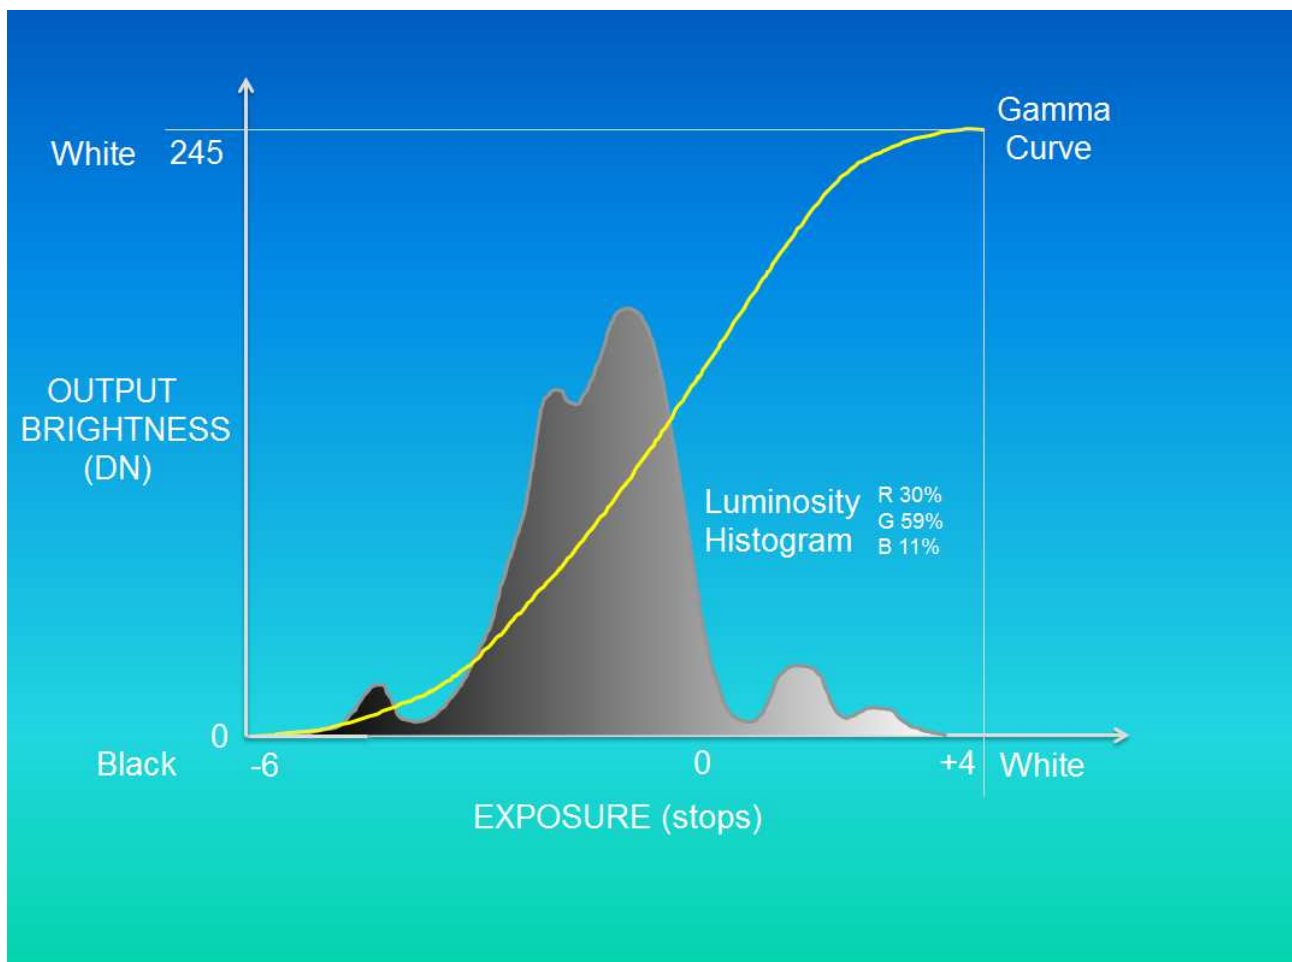


Fig. 6 This image shows an unnatural looking blue sky, the result of making a global saturation and brightness adjustment. The standard blue sky has R = 135, G = 206 and B = 235 DN.

5. A little less noise please!

Noise is fatal for many images. To understand how to manage it we need to go back to the moment the image was captured. We have already mentioned that it is an advantage to over-expose the bird, provided no highlights are fatally blown out. The brightness is then reduced in post-processing to achieve optimal exposure. A side-benefit of this process is noise reduction (NR). There are a number of noise sources, some of which, like read-out noise, are independent of the exposure. This means that the amount of noise recorded will be the same whether the image is under- or over-exposed. The noise will be proportionately greater in an under-exposed image compared with an over-exposed image. In technical jargon we say that the over-exposed image has a higher signal-to-noise ratio (SNR) than the under-exposed image. When the brightness slider is used to pull back the over-exposed image it retains the more favourable SNR. If you have an underexposed image (lower SNR) and use the slider to pull up the brightness, then you will find the image still has a low SNR but because the brightness has increased, the noise will be even more apparent.

There is also something else going on which helps to reduce all kinds of noise in the dark parts of the image. This is where noise appears to be worst and when the brightness slider is pulled back, the dark parts of the histogram are pushed to the left under the heel of the gamma curve (see Fig. 3 – copied from Part 1, below) where both the brightness and contrast are reduced more than for the mid-tones. It is like sweeping the noise in the dark parts under the carpet. Of course, the opposite happens if you have an underexposed image and pull up the brightness; the noise comes out from under the heel of the curve, becomes brighter and picks up contrast.



Recording an image that is a bit over-exposed is a good place to start but there is one more thing we can do in-camera. It may seem counter-productive but it is usually an advantage to turn off the camera's internal NR if you have the option. This is because in-camera NR reduces the sharpness of the image, something we wish to avoid and there are more effective ways to deal with noise in post-processing. You may also find that when you download your image to a RAW converter such as Canon DPP, the software assigns default values for NR. There will be two kinds of noise; luminance noise and chrominance (colour) noise. Luminance noise is the more troublesome and is the kind of noise needing special attention in post-processing so the luminance NR should be set to zero, otherwise it will reduce sharpness. We will deal with the luminance noise later. On the other hand, chrominance noise, which appears as RGB speckle due to errors in colour digitizing (demosaic process), is very effectively reduced using the options available in RAW converters. If you are using Canon DPP, the software will recognize your camera body type and ISO setting and provide an effective default setting for chrominance NR. If you are using a third party RAW converter such as Lightroom or ACR you will need to use your judgement to adjust the chrominance (colour) NR slider.

There is one more NR 'free kick' available if we are reducing the size of an image for web display. The downsizing process involves a mathematical resampling of the pixels in the image to calculate new values of brightness based upon surrounding values. It is a bit like averaging a bunch of neighbouring pixels to arrive at a new value but it is more sophisticated than this and produces a more faithful version of the original image than one can achieve with simple averaging. Interpolation using a bicubic function is used instead of an average so whenever you downsize an image you will see several options for resampling; click on the 'bicubic' option (do not use bicubic sharper or bicubic smoother). During the process, good pixels and noisy pixels are used and provided there are more good pixels the end result is a degree of NR. I call this a free kick because the downsizing was necessary for web display and as a by-product we got some NR. It is worth keeping this in mind when judging how much NR is required in post-processing; remember that you will get a little bit of extra NR in the downsize and therefore you need a little bit less in the post-processing.

So, what NR are we going to do in post-processing? This can be confusing for the inexperienced because there are many options from which to choose including third party plug-ins such as Neat Image. Some are very effective but you need to be aware that, to some extent, they all reduce sharpness, so there will be a trade-off between NR and sharpness. A good strategy is to only use NR on those parts of the image where it is really needed, this is called selective NR. The parts that are most in need of selective NR are the dark parts and out of focus background and foreground. In these parts of the image, noise is most obvious and where we should focus our NR effort. In general, it is hard to see noise in those parts that are in sharp focus and rich in detail such as we find on most birds. This means that in most cases it is unnecessary to apply any NR to the bird unless you are using a very high ISO. The first step is to select the bird and its perch using one of the selection tools and then switch to the inverse so that the background and foreground are selected. The selection tool should be hard-edged with feather = 0. The inverse selection will protect the bird and perch from any adjustments we make to the background/foreground. Remember to save the selection because it will be needed later.

Most post-processing software will offer a number of NR tools. For background/foreground NR they are all usually effective. One of the simplest is the median filter with radius set to 1–5 pixels; it is very effective on most forms of noise and preserves detail reasonably well. Advanced NR software is available which automatically targets specific areas of noise in the image and optimizes the amount of NR that is applied depending upon the brightness and amount of noise present. You can do this yourself by manually selecting noisy parts of the image with a soft-edge tool and applying just the right amount of NR. Sometimes, parts of the bird will have unacceptable noise and this too can be managed with selective area NR. Usually I find this unnecessary and just one round of NR on the entire background/foreground is all that is required. The process I have outlined is the simplest kind of selective NR but you will find it is surprisingly effective in most cases.

An extreme example of NR is illustrated in figs 7(a–c), a black and white bird captured using my Canon 5DIII camera with the sensitivity set to ISO 12,800. I use this image to test NR and detail recovery techniques. Fig. 7(a) shows a zoomed in view of the noise in the bird's head area and 7(b) shows the same area after NR using the Neat Image plug-in and finished with a median filter on the background. The full frame is shown in fig. 7(c).



Fig. 7(a). Part of an image recorded using a Canon 5DIII camera with the sensitivity set to ISO 12,800.



Fig. 7(b). The same area following noise reduction using Neat Image and after cleaning up some artefacts in the background with a median filter.



Fig. 7(c). The full image after finishing.

6. Sharpening know-how

Too much or too little sharpening is another common deadly sin. Some sharpening is always required when processing images for our gallery but many members do not know at which stage in the workflow to apply sharpening and how much to apply. Also, there are usually several sharpening options and the choice can be perplexing.

Let us start back at the point of image capture. You are probably not aware that your camera will apply some sharpening that is out of your control during image capture. In addition, most cameras also allow the photographer to set some in-camera sharpening. If you are shooting RAW you will probably apply some in-camera sharpening so that the image you see on the camera LCD is reasonably sharp. When your image is downloaded to your RAW converter, the in-camera sharpening will either be stripped off or, if you use Canon DPP, it will be retained as a default sharpening value which you can accept or adjust to your liking. Some sharpening is essential during RAW conversion because the colour demosaic process results in a mild form of blur. The amount of sharpening applied during RAW conversion, known as capture sharpening, should be just enough to notice when viewing the image at normal size. You will probably have some sharpening options including the popular Un-Sharp Mask (USM). The USM threshold parameter is used for noise reduction and will reduce sharpness so it should be set to zero. USM also introduces undesirable artefacts and is not recommended for capture sharpening. If you use Canon DPP, then the single parameter 'Sharpness' provides optimal capture sharpening when set to 3 on the 0 to 10 scale. DPP also offers the Digital Lens Optimizer (DLO), a unique capability to perform a mathematical operation known as deconvolution which partially removes blurring due to residual lens aberrations, the optical low pass filter and diffraction. It is like a treble boost in a hi-fi system; the end result is a sharper image but it will also increase the noise. The DLO slider should be set to a value of about 50.

If you are shooting JPEGs you should not use any additional in-camera sharpening as the manufacturer has already programmed the camera to apply the optimal amount of sharpening for your camera model. When you download for post-processing you should not apply any sharpening until the image has been downsized for our gallery.

Downsizing and output sharpening are the final stages of the post-processing workflow. We mentioned earlier that downsizing results in some beneficial NR but it also results in a loss of sharpness. This is because the bicubic resampling introduces a mild blur like Gaussian blur. One of the downsize options available in Photoshop, and some other processing software, is 'bicubic sharper'. This option uses bicubic resampling followed by automatic sharpening but it is a global adjustment and the amount of sharpening is no longer in our control. For these reasons, it is not a good idea to use this option. Once again, there are usually a number of options for sharpening the downsized image, including USM, but it should be used carefully and with the threshold set to zero, unless noise is a problem. In Photoshop, there is another option in the Adjust Sharpness dialogue box called 'remove Gaussian blur'. It uses deconvolution to remove most of the blur resulting from the downsizing. I say most of the blur because the downsize blur is not exactly Gaussian but the deconvolution does a good job nevertheless. I prefer this option instead of USM because the resulting image has fewer artefacts and is therefore a more faithful copy of the original image. The optimum parameters to remove Gaussian blur are radius 0.3–0.5 pixels, amount 30–75% and you should check the 'more refined' box. Fig. 8 shows a screen grab of the Adjust Sharpness dialogue box as it appears using my ancient version of PS Elements.

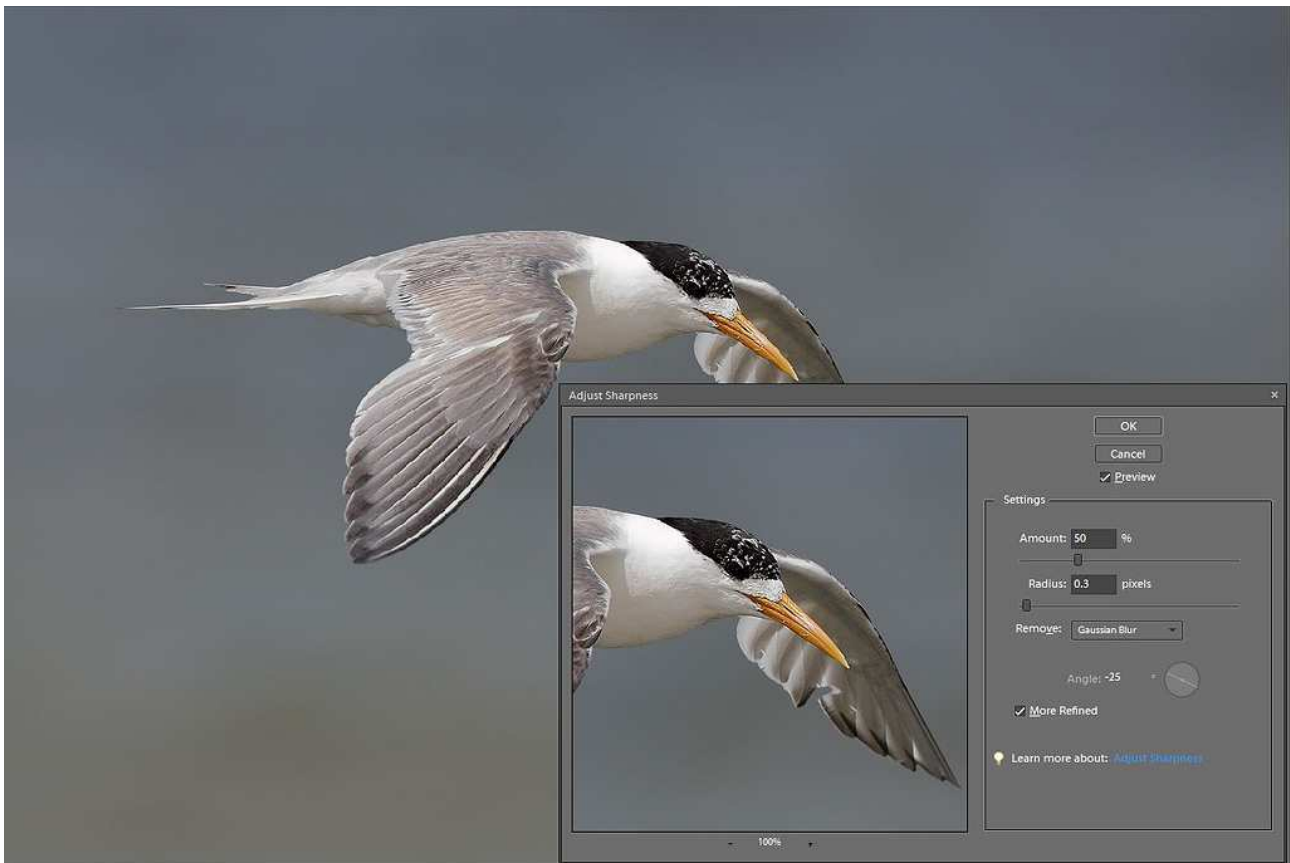


Fig. 8. A screen grab of the Adjust Sharpness dialogue box in PS Elements 7. The remove Gaussian blur option should be used for output sharpening.

Finally, you should only apply the sharpening to the bird. This is where you need to reload the bird selection made earlier while managing noise. You can also apply sharpening to selected areas of the bird if needed. After downsizing and output sharpening it only remains to convert and save the image as a JPEG of suitable file size for the gallery. In Photoshop you do this using the Save for Web dialogue box and remember to check the Colour Profile box so that the image will display with correct colours on everyone's computer monitor.

In summary, if you are shooting RAW, you need to do some mild capture sharpening before RAW conversion and output sharpening after downsizing. If you are shooting JPEGs, you should sharpen once after downsizing (output sharpening). There should be no need to sharpen at any other stage in your workflow.

7. Composition is really important

By good composition I mean arranging the picture elements so that the overall effect is pleasing to the eye or captures an interesting story. It is a really important quality that makes a great deal of difference to the presentation of one's work. Unfortunately, my survey revealed that about 60% of pictures in our gallery have compromised composition but, fortunately, it is usually one of the easiest problems to fix. Very often it is just a question of cropping; choosing landscape or portrait, the right frame proportions, putting the bird in the right place and giving it some 'space'. There are well-known rules of composition that are helpful and there is a wealth of experience of what works and what does not work in bird photography. We do not have the time or space to discuss the subject properly but, instead, recommend members take a look at some online bird

photography tutorials such as the presentation by British pro, Mike Atkinson <http://mikeatkinson.net/Tutorial-6-Basic-Composition.htm>.

Ideally, good composition starts in the camera and finishes in the computer where final adjustments can be made. In the field, one should try to position the bird in the frame somewhere near the optimal position depending upon the pose of the bird, interesting action or unusual lighting. To do this one needs to focus and then recompose the image or use an off-centre AF point. Often this is not possible, for example with flight shots, where one must put the bird in the centre of the frame under the central AF point and recompose later by cropping the image in the computer. Generally, the 'rule of thirds' is a good guide to subject placement but there are exceptions as discussed by Mike Atkinson and others. Fig. 9 is an example of the use of the rule of thirds.




Fig. 9. An illustration of the use of the 'rule of thirds' for subject placement in the frame.

It is also an advantage to shoot at eye level with the bird. This will often mean getting down low or even into water. The award-winning Canadian photographer, Gail Bisson, wears chest-high waders and sets up her camera just above the water on a tripod. A low point of view has the advantage of causing most of the foreground and background to be out of focus and minimizes harsh shadows on the ground. The right head angle and eye-contact with the bird are highly desirable as they can provide a powerful focus of attention for the viewer. Avoid distracting elements in the foreground, including other birds, and try to compose the scene with a background showing something of the environment, otherwise try to choose a smooth uncluttered background. To some extent, the background can be manipulated in post-processing to make it smoother, lighter or darker, helping to make the bird stand out more or 'pop'. Distracting elements can often be removed using the clone tool, and if you need more space for the bird you can use the Canvas Extension dialogue box to extend the edge of the original frame outwards. This is particularly easy to do with simple backgrounds like sky and out of focus smooth backgrounds.

One of the secrets of good composition is keeping the scene simple with just a few picture elements that focus attention on the bird. One of the simplest image designs has three elements; the bird, a perch, and a smooth out of focus background. This works very well and is a popular choice but beware of manipulating the background too much as it can look over-worked, even fake. Try to retain some natural variability in the background but tone down any distracting areas such as bright patches of sky. You should also try to shoot with backgrounds having colours, textures and patterns that are in harmony with the bird and try to avoid man-made perches and built environments unless they are part of the photographic narrative.

My final point addresses what I see as the most common compositional shortcoming on our gallery; not giving the bird enough space, in other words, cropping too tightly and, in the extreme, trimming feet, tail and wings. When one captures a nice image with lots of feather detail there is a natural inclination to show off the technical achievement by making the bird as large as possible in the frame. Tight cropping is usually fatal and spoils the presentation of what might otherwise be a wonderful image. A related matter is subject placement in the frame. Within the frame, give the bird some space to look, move or fly into. This can be up to half the width of the frame. Space in the right place helps to provide scale and context to an image.

To learn more about composition, take every opportunity to study award-winning photos, look at the work of professional bird photographers on galleries such as BPN, and join in the discussion of images on online forums.

SUMMARY

To conclude this talk I have prepared the following check-list that I recommend members use before posting images to our gallery. If you tick all the boxes you can be confident that your image will be well received and admired by most members. It will also better serve the main objective of the gallery; to record high quality images of every Australian bird species, male and female, in all plumages and stages of life. Let us all try harder to lift the standard of our work to meet these personal and collegiate objectives.

1. Check that the bird is in focus with acceptable depth of field.
2. Is the bird well exposed?
3. Are there any blown highlights and have you extracted the detail in the light and dark parts of the bird?
4. Is there a colour cast, do the colours of the bird and any sky or water look right, and is the colour over-saturated?
5. Is there unacceptable noise in the image?
6. Has the image been properly sharpened?
7. Have you achieved an aesthetically pleasing composition or interesting photographic narrative?

ACKNOWLEDGEMENTS

I would like to thank the DPITB organizing committee for inviting me to make this presentation at the Newcastle event. John Stirling kindly made time to review and critique the text, while Jill Wilson filtered out little literary indiscretions.