Slide1
This presentation will be looking at ways of making the most of the bits and pixels in the camera output picture file, in particular, ways of correcting exposure errors in the picture. This morning Tony was looking at RAW converters, which are analogous to dealing with the film negative, whereas this presentation will be dealing with the rendered eight bit JPEG or TIFF file which is more like the output print. All being well the output file produces the desired picture or print for the intended application, but where this is not the case we will be looking at what can be done to produce the desired picture. I will also be looking at interpolation methods and how far it is reasonable to go with print sizes.

Slide 2
Primarily we will be looking at digital camera images for display or for making prints, and looking at some of the common problems that result in exposure errors and ways of correcting them. Have a quick look at auto-fix algorithms, and then summarise.

Slide 3
I looked through my last few years of pictures and could not really find that many problem pictures but these are some examples with exposure errors, where the camera control has not achieved the desired exposure, the camera can be misled by too much sky or the light colour of sand on a beach or in snow scenes. Cameras now-a-days have many exposure modes to help in such circumstances, but sometimes the picture is taken before one realises that there was a better exposure programme and then of course the moment of the picture has passed.

Under exposed images generally have more scope for correction, most of the scene information is just hidden in the shadows, unfortunately, if any data is clipped in either the highlights or shadows then it has been lost, unless you can get back to the RAW file.

Over exposed images are more of a problem, it is likely that more image data will be clipped along with the highlights. Close flashed images can also be difficult as the majority of the data is in the highlights.

Slide 4
These are the sort of auto fix processes that might be used when a picture is taken, in the camera, the auto white and auto exposure options, make it easy to take pictures in varied situations. Compared to a film camera/film printer system the digital camera has to perform both the image capture function as well as the print processing function, in order to produce an output file suitable for display or printing. So… it is not surprising that it occasionally gets it wrong. In comparison to the camera auto functions, the auto functions in photo-processing software packages like PhotoShop, are more generic and can give good results but their parameters are set for the typical picture and may not be optimum for the picture you are trying to correct.

We will look first at how levels is used in photo processing software to improve the appearance of pictures.

Slide 5

This slide shows a slightly dull looking picture of a lion, the green separation and the green histogram. The histogram shows the number of pixels at each code value, starting with a pixel value of zero at the lhs through to a pixel value of 255 at the rhs. The higher the graph the higher, the number of pixels at a particular code value.

In order to obtain a bright and colourful picture we need to use the full range of code values available, otherwise the picture will tend to look dull, as in this case, or too dark or too light, depending on the actual range used.

PhotoShop auto-contrast and auto-levels adjusts the gain of the picture data to fill the full range of code values. The gain is adjusted to clip a fixed proportion of the pixels at zero and 255, in practice though, the optimum clip value will depend on the picture.

The second histogram has the lower input slider set to 20 meaning that all input code values at 20 or less will be set to zero. At the upper end the slider is set to 200 meaning that all input code values at 200 and above will be set to 255. The bottom histogram shows the result, where the data between the clip values has been stretched to fill the full range, but there are only 180 (200-20) code values for the positions 1 to 254 leaving 74 gaps.

The centre slider can be used vary the tone scale giving a two slope function, while maintaining the high light and shadow points.

This is perhaps a good time to apologise to the film members of the audience, I was brought up in the television industry where words such as contrast and gamma do not mean quite the same thing as in the film community but I trust you will bear with me in the use of these and other words.
The lower pictures now show the effect of the change of levels, and lhs picture is certainly brighter and more colourful than the one at the top. Setting of the sliders to 20 and 200 has clipped around 0.5% of the pixels in the shadows and the highlights and it is easy to see the clipped pixels in the highlights, on the lion, in the corrected green channel. 0.5% is the same clip setting that Photoshop uses in its auto levels and auto contrast function.

Slide 6

First Auto Contrast, if there are already a proportion of pixels clipped then carrying out any auto levels function will not change anything. PhotoShop auto-contrast averages the red, green and blue channels first and then works out the code value to give the clip level of 0.5%, that code value is then used in each colour, thus applying the same gain increase to each colour. In this example there is only a very small change to the picture. The average of the three channels already exceeds the proportion of clipped pixels at zero and the algorithm has only slightly increased the clipping in the highlights. The histogram on the right has only a single gap visible, confirming that there has only been a small change in gain.

But in Photoshop auto-levels, the same percentage clip level is used in all three colours, rather than the same code value in each colour. This will produce a more neutrally balanced picture, which is fine for under exposed pictures and pictures with an unintentional colour cast.

Slide 7

But for highly coloured scenes such as sunsets and sunrises it does the wrong thing. This picture shows what happens if auto-levels is now applied to my picture of the rush hour, over Milton Keynes; it has changed from early morning to mid morning. The histograms on the rhs show that the gain has been increased to clip a the proportion of pixels at 255 in each colour. As can be seen, the gain is highest in blue where the sunrise colour balance has under-exposed the blue channel by about two stops. The green channel is about one stop under exposed and there is only a very small change in red.

Auto contrast will cause the least change to the picture and is the safest to use if you do not want a colour balance change, but even so, using the levels function to correct an exposure error is not right.

Slide 8

When the gain is adjusted in 'levels', the exposure error is being corrected by changing the gain of the data in display space. To accurately correct exposure errors the gain change should be carried out in the same space in which the error occurred, in this case, scene space. In order to do this the camera data has to be first unbuilt, to take it back to scene space, the exposure error corrected and the data passed through the camera transfer function, back to display space. But... there are two main problems with this approach, the first is knowing the camera transfer function in order to unbuild it, and the second is the use of eight bit data. For the first problem, the camera transfer function, there is a standard. For the second problem we either work with sixteen bit data or apply a non-linear correction to the data after the gain adjustment. Which may well be, what the majority of you would do, after changing levels, in order to obtain an acceptable tone scale.

Slide 9
This is a block diagram of the exposure correction process, where the data is first unbuilt to scene space, with a characteristic similar to the curve shown. Once in scene space, any errors to the exposure can be corrected by changing the gain. After which the data can be passed back through the camera transfer function to return it to display space.

Slide 10
ISO 14524 describes methods for measuring the opto-electronic conversion function of digital still cameras and specifies methods for focal plane measurements which are analogous to film methods, as well as whole camera measurements.
This standard uses a test chart, comprising a series of twelve calibrated patches in a circle, and specifies both the test chart and the methods of illumination for reflection and transmission charts. The standard also specifies that the results are to be displayed as the digital output vs. scene log luminance for the whole camera measurements, similar to the way film characteristics are shown. For the purpose of this presentation where we are trying to correct exposure errors it is more useful for the results to be displayed as digital output vs. relative scene luminance.

In order to display the pictures from a digital camera correctly on a soft display the file format usually stores the data in sRGB colour space, although other colour spaces can be specified, I will assume sRGB.

Slide 11
This standard specifies the reference display and observer conditions as well as the viewing conditions for the default sRGB colour space. It also defines the encoding transformations.
The sRGB characteristic is essentially an inverse power law but with a limit to the maximum gain, in the blacks, of about thirteen. This tells you that to accurately obtain eight bit sRGB data you need to start with at least twelve bits of data from the camera A-D converter.
In order to first unbuild the camera data the inverse characteristic, as shown in the lower curve and the lower equation, has to be applied.

Slide 12
I am afraid this is the limit of my test picture ability, a couple of old muppets on a bench and a home made colour chart that bears some resemblance to a MacBeth chart, which I am sure Mike will describe and show much more accurately. Nevertheless this picture will suffice to demonstrate the accuracy of the process, by having a number of reference colours, with which we can make any comparisons.

This is a do nothing process with the normally exposed test picture, or more accurately what my camera auto-exposure thought was the correct exposure. Starting off on the lhs, with the camera original, the data is first passed through the camera unbuild to bring it

back to scene space, in which any exposure errors can be corrected, the corrected image data is then passed through the camera transfer function to get back to an image we can display or print.
The green histograms associated with each picture have now been added, by using a single colour histogram the gaps in the code values are more visible. The vertical scaling is set by Photoshop and is not always the same, so this has to be taken into consideration when comparing some of the histograms.

The centre picture has been transformed to scene space by undoing the camera transfer function, the histogram already shows gaps at the higher levels due to the steepness of the inverse transfer function, while at the lower end the characteristic has compressed the data, leading to some data being discarded, but of course there are still the same number of pixels, so the pixels that occupied the upper code value gaps have now dropped into some of the lower code value bins, considerably increasing their numbers. In this case there is no gain adjustment, but when the data is passed back through the camera transfer function, to get back to our original image, there are even more gaps in the histogram, this time at the lower levels, where the discarded data would have been and once data has been discarded it is not possible to regain it. The gaps in the final histogram are showing that a considerable number of code values have been lost.

Slide 13
In a typical picture a few missing levels well distributed would be difficult to see, particularly if they are at different spatial positions in each colour, but a cluster of missing levels or large gaps will show up as contouring, or posterisation, particularly in smooth areas. But in a typical picture, textures or camera noise will help to mask the contouring making it difficult to see.

If there are not too many missing code the easiest way to regenerate them is simply to save the file as a JPEG. This may go a little against the grain for some but it will blur the edges of contouring with the JPEG blocks. The file can then be resaved as a TIFF to minimise any further degradation but if the processing has produced a large number of missing levels then the picture is probably not of the highest quality so saving it in JPEG will not noticeably degrade it any further, and will help reduce the visibility of contouring.

Slide 14
This slide shows three blocks of greyscale each with three rows, in the top row the increment is one level between steps, equivalent to eight bit data, the second row has an increment of two levels between steps, equivalent to seven bits and the bottom row has an increment of four levels equivalent to six bits. The steps are certainly all visible on the bottom row and on the middle row they are just about visible for all but the extreme black and white ends.
But if things are made a bit more realistic, by adding some noise, as shown in the lower set of greyscales, then it is more difficult to see the steps in the middle seven bit row, although steps are still visible in most of the lower row, so it looks as though in a typical picture it would be difficult to see the contouring in a seven bit picture, which gives us

scope for up to x2 in gain correction without any visible loss of quality and without the need to use the JPEG save trick.

These blocks have all three colours changing step together, in practice the three colours are unlikely to change together, making the contouring even less visible.

Slide 15
Looking back at the unbuild/rebuild of the normal file then after the JPEG save the green histogram looks much smoother and closer to the original. The number of colours in original is 312k, but it is down to only 220k in re-built TIFF picture, due to the missing code values. But after the JPEG save, the number of colours has risen to 330k, more than the original, but this does not mean the original colours have been regenerated.

Slide 16
We just about got away with unbuild/rebuild with the normal, without doing any gain adjustment, but if the same process is carried out with the -2 stops picture the limitations of using eight bit data becomes very apparent. After the gain correction the gaps are very visible, but after gamma correction the gaps have become very large and unacceptable, particularly at the black end.

Slide 17
If the JPEG save trick again is used again at the same compression setting of 10, you can still see where the gaps were; the quality setting needs to be reduced to 7, with the associated loss in picture quality, to obtain a smooth histogram. To maintain our processed picture quality we need a different approach, in order to hold onto as many pixel values as we can.

Slide 18
If the gain adjustment is carried out first then all the pixel values are kept, with the exception of any clipped pixels in the shadows or highlights. Because the gain adjustment has been carried out in sRGB space a small non-linear correction, to boost the relative levels of the dark pixels, needs to be applied. Knowing the camera transfer function, the non-linear correction for different levels of gain boost, can be calculated.

Alternatively an exposure series of pictures can be taken and correction curves generated, using the ‘curves’ function, to match the exposure series to the ‘normally’ exposed reference picture. The correction curves can be saved for each gain setting and called up to correct any other pictures, from the same camera, where a gain adjustment has been made.

Slide 19
Looking at the 1 stop under picture, the histogram data occupies about three quarters of the range confirming that it is about -1 stop. The camera transfer is very approximately an inverse square law and the square root of 0.5 is about 0.7.

First the gain is increased to give the same clip level as in the 'Normal' that is being used as a reference, this results in some gaps in the histogram. After the gain change the tone scale has to be corrected with a small non-linear correction in order to match the processed -1 stop, to the normal.

Slide 20

This slide compares the 'normal' with the two methods of exposure correction for the 1 stop under exposed picture and confirms that the two methods give similar pictorial results and that they are a reasonable match to the 'Normal' but that the histograms are very different. The gain/gamma method has far fewer gaps in the histogram compared to the unbuild/rebuild method and would give a much more acceptable print.

Slide 21

With the one stop under exposed picture we can do a bit better and maintain some of the highlight information that is missing from the 'Normally' exposed picture. Rather than hard clipping the highlight pixels to 255 a soft clip characteristic, as shown, will compress the highlight information while still maintaining some highlight detail. There will now be a slight difference between the normal and the –one stop picture in the highlights, depending on the level of soft clipping.

The curve shows that the soft clip has reduced the maximum code value, the gain could be increased at the same time using the 'curves' function but then it would be difficult to keep track of the gain change. The value of the gain change is needed from the 'levels' function in order to call up the appropriate non linear correction curve.

Slide 22

Applying the same process to the -2 stops exposure, there is a little more curvature to the non-linear correction but even so the gaps in the histogram are fairly well distributed and quite acceptable.

Slide 23

This shows the comparison between the two methods, again there is a reasonable pictorial match between the results and a large improvement in the histogram for the second method.

If the file is saved as a JPEG with a setting of 10 in PhotoShop, the gain/gamma method shows a fairly smooth histogram, whereas the unbuild/rebuild method still shows holes rather than gaps in the histogram. The histogram can be smoothed out further by using a lower quality setting but it is a compromise between JPEG artefacts and contouring to obtain the best final print.

Slide 24

Looking now at the 1 stop over exposed picture, things are a little more difficult.

For over exposed pictures, the problem is that there is generally a large number of pixels clipped in the highlights and having been clipped there is no way to regain the data, sometimes when only one channel is clipped it is possible to copy the highlight data from the unclipped channels into the clipped channel.

Starting at the lhs, the original is first unbuilt using the inverse of the camera transfer function.
In the gain correction stage, because the picture has been overexposed, it is necessary to reduce the gain, but because the highlights have been clipped, there are no pixels available to occupy the upper code values, so in order to make this stage work, with eight bit data we have to cheat a little and maintain the already clipped pixels at 255, and this explains the spread out histogram above the half way point. If this is not done the output picture will have no values above about 180. The dotted line shows the halving of the gain for the main body of the picture data. After the gain correction the picture data is passed back through the camera transfer function.

Slide 25
In the gain/gamma method, in adjusting the gain there were more than enough clipped highlight pixels but at the shadow end the soft clip method can be used again to prevent blocked out shadows. It not essential and the effect will not be as useful as using the soft clip in the highlights of the 1 stop under exposed, but it is still better than hard clipping the extreme shadow data. The histogram shows that there was only a small gain change. Because the picture was over exposed, then the non-linear correction curves in the opposite direction, to the under exposed correction. The best that can be done is to correct what you have and accept that there will be large areas of the picture clipped in all colours and some of the lighter colours will remain incorrect because one or two of their channels are already clipped.

Slide 26
When the two methods are compared the histograms show that the gain/correction method still has far fewer gaps than the unbuild/rebuild method. Despite the cheat in the unbuild/rebuild method the pictorial results are very similar, and quite a few of the colours are a reasonable match to the reference 'Normal', depending on how tolerant you are. Most of the colours are looking more saturated because one or more of their channels were clipped. The trouble is that there is no way of telling if the clipped pixels are only just clipped or are severely clipped, and this removes the correct relationship between the channels.

Slide 27
Looking at the corrected pictures vs. auto-levels there is a big improvement in the tone scale compared to Auto-levels alone, which is not surprising since the levels change is the first stage of the gain/gamma method. The correction curves used in these pictures can be applied to any pictures from the same camera requiring the same gain change. In these examples the correction has been applied equally to all colours but the correction could be applies separately in each channel to correct an unwanted colour cast.

One interesting side effect is that the sharpening level looks different on the corrected images. This is the sharpening that is applied in the camera and it is generally applied in the output display space. The exposure corrections are effectively applied in scene space so the relationship between the tone scale and the enhancement has been changed.

Slide 28

Applying this method to a couple of the problem pictures, this picture may have lost its 'atmosphere' but ignoring that, the resultant tone scale shows a big improvement over just using levels to correct the picture exposure error.

Slide 29

This is a combined beach and snow scene and the corrected picture shows an improvement in the tone scale in particular to the background huts.

Slide 30

The camera original histogram shows the typical bi-modal distribution of this backlit picture, where the underexposed foreground is to the lhs of the dotted line, and the overexposed background is to the rhs of the dotted line. The dotted line is about half way along indicating that the foreground is about two stops underexposed. The foreground can be selected, the gain boosted and the –two stops non-linear correction applied to give the corrected centre picture. The top right hand histogram shows that the corrected under exposed data is now spanning the full range of code values, albeit with gaps. After a JPEG save the gaps have been filled giving a smooth histogram. Since the original picture already has clipped pixels, applying PhotoShop auto levels or auto contrast will make little difference.

Slide 31

In conclusion to this part of the presentation, to implement the gain gamma method of exposure correction the easiest way is to take a series of pictures with different exposure, using the camera manual settings where possible. If that is not possible then use different values of the camera EV setting, failing all that you could always resort to using neutral densities in front of the lens after locking in the exposure setting.

Secondly use soft clips, before levels, to decrease the dynamic range of the input picture, and maintain highlight and shadow detail.

Having produced the exposure series, they can be processed in PhotoShop, or other photo-processing software that has the levels and curves function. The correction curve, that matches the 'Normal' for the different exposures is stored with a file name you can recognise.

If your processing and has generated large gaps in the histogram, then the JPEG save may help in overcoming any contouring.

Slide 32

Having optimised the tone scale, how large can the picture now be printed. We will be looking at various interpolation methods for enlargements and the sort of results you might get. The maximum picture sizes you might achieve and what you should aim for in preparing your picture.

Slide 33
There are two main application and they each have slightly different requirements. The first application is a whole picture enlargement, to make a poster or a large print to hang on the wall.
…and the second is where you want to make a selective enlargement, probably to make a normal size print or enlargement.

Slide 34
Looking at the various interpolation methods the top four are readily available in most photo-processing packages, the Photo-Zoom S-spline is a stand alone application. I have a free trial version that imprints a watermark on the pictures, as you will see later. The Fractal software I have is also a free download that lasts thirty days and is a plugin for Photoshop.

Pixel replicate or nearest neighbour is not really an interpolation method it simply takes the value of the nearest pixel and assigns it to the new pixel. It can be used to give a pixellated effect.

Bilinear as its name suggests is two directional and calculates the value of the new pixels as a proportion of the linear distance to the four surrounding pixels, above, below, to the left and to the right.

Bicubic uses the sixteen surrounding pixels to determine the new pixel and uses a cubic function, of the distance to the new pixel, to calculate its value. There are softer and sharper options available in the current PhotoShop offerings. In the examples I have used the standard Bicubic setting plus some unsharp masking to match the sharpness to the other methods

Stair step is a method of using bicubic interpolation where the interpolation is carried out in a large number of small increments rather than in one step. It can sometimes offer an improvement over single stage bicubic but there are schools of thought that suggest every stage of interpolation causes some degradation and so one should minimise the number of interpolation steps.

S-Spline tries to fit an equation to the surrounding pixels, and then uses that equation to determine the value of the new pixel, generating a new equation for each new pixel, unlike bicubic and bilinear that use the same equation for each new pixels, albeit with different coefficients.

Fractal enlargement samples from a larger number of pixels than bicubic and seeks repeated patterns in decreasing sizes of pixel blocks, relying on the similarities that recur at different scales in nature. Fractal patterns can be scaled to any size without loss of detail, though the image quality will depend on that of the fractal analysis.

In the examples that follow I will only be comparing Bicubic, S-spline and Fractals

Slide 35
For those not familiar with the PhotoZoom, it is a stand alone programme and offers a wide range of customisation in implementing the interpolation function, such as unsharp masking, edge, detail and sharpness boost levels as well as addition of artificial detail, as shown is the enlarged shot.

Slide 36
The interpolation methods cover far more than those available in photoShop, in addition a number of preset settings can be called up as well as the users own settings.

For the examples in this presentation I have kept with S-Spline XL and Photo-detailed. With so many variations available, this programme could keep one amused for a long time.

Slide 37
This is the genuine Fractals screen shot, this is a more dedicated programme and is implemented as a plugin and will only work with PhotoShop CS2 or 3 and elements 4 or 5. It also needs to have dotNET framework 2.0 installed, which can be downloaded from the Microsoft web site.

The control layout does not seem to have been designed with visibility in mind, with grey text on a grey background.

Slide 38
But to help, I have overlaid the additional controls which are texture, sharpening and film grain. Again for the examples in this presentation I have used these default settings.

Slide 39
The picture in this example is relatively low noise with little texture, the slide shows two levels of zoom x2 and x4. At the x2 there is little to choose with the s-spline looking the sharper and smoother. At x4 the s-spline is definitely sharper followed by the Fratal version and then bicubic.

Slide 40
This slide shows two further picture enlargement x6 and x8. You probably need to look at the complete picture before coming to a conclusion, but for this part the S-spline method looks the sharper in both cases followed by the Fractal method and bicubic. It looks as though the S-spline method could go on further, but while it is sharp there is no further data, whereas, in theory, the Fractal method should be generating new detail data.

Slide 41
This picture is the opposite of the previous one, it has a slightly higher noise level but has fine texture in the cats fur and whiskers. For this slide and the next couple, a number of pictures were taken with variable optical zoom settings with the idea of comparing digital zoom of the various interpolation algorithms with the camera optical zoom. The camera originals show the two optical zoom settings and the other pictures show how the interpolation methods compare with each other and with the optical zoom. The optical zoom on my camera is not scaled so I had to guesstimate the point of zoom of each set of pictures, hence 175%. At this setting the interpolation methods give similar results and are not far away from the optical, unfortunately my cat has moved between shots so you will have to use your imagination a bit.

Slide 42
X4 zoom still has good definition of the whiskers and fur of the cat albeit somewhat coarser than the optical. There is little difference between the methods, although the Fractal version may have slightly finer features.

Slide 43
X6 zoom this is the limit of my camera zoom, but the interpolated image has lost most of the definition in the whiskers of the cat and there is little difference between the methods but the fractal method still has the slightly finer features.

Slide 44
The resolution limit of the eye is around 5 cycles/millimetre at a viewing distance of 30 cm so, for print sizes up to about 10x8, this gives us a resolution of 10 pixels per millimetre or 250 pixels per inch, you could print at a slightly higher resolution to overcome any losses in the system. For this chart we will keep with 250 ppi for quality prints and 150 for just acceptable prints and posters, with their greater viewing distance.
The ppi assumes an image source with three sharp colours, which is true of a three sensor camera or one using the Foveon sensor or a tri-linear scanner. The vast majority of digital cameras use a single sensor, where half the pixels are green and a quarter each for red and blue, which means we should really halve the picture size but that may be a little harsh, the manufacturers use sophisticated adaptive algorithms to regenerate the missing colour pixels, a reduction of picture size to 75% rather than 50% may be more realistic.

Slide 45
This is the same chart but in inches, so for a 6 M pixel single chip camera the high quality print size is 8.5 x 6.5 inches and the x1 poster is 14 x 11 inches.

Slide 46
To achieve the best results from the camera you should use the highest settings, for both pixel number and file quality.

Use the sharpness setting that gives the sharpest picture with no overshoots, this may be different with different pictures, a high contrast scene would require a lower sharpness setting than a scene with low contrast texture.

If the ISO setting is available then in general the lower the setting the lower the noise, but there may be a point in the ISO setting, below which, there is no further reduction in the noise.

If the shutter times are getting a bit slow, say less than 1/60, then a tripod is essential, or use a flash or even both

Slide 47
So how far can you go in picture enlargements, for a typical picture up to x4 enlargement for high quality prints or posters, using any of the interpolation methods, for a typical 6 Mpixel camera the corresponding print size is 32 x24 inches and for the poster up to 56 x 44 inches. But with a good low noise picture with little texture, using S-spline you could go up to over 100 x 80 inches, you have just got to go out and buy the printer.