Chapter 6

Colour FPN Correction

6.1 Introduction

In the previous chapter, the nature of colour, its capture, perception and interaction with imager sensors was detailed leading to the adoption of a perceptually uniform colour error metric. However, this discussion did not include the effects of the three main sources of error - model error, residual FPN and quantisation noise along an imaging chain - on colour accuracy. It assumed a perfect response model and an ideal manufacturing process without due consideration for device mismatches causing fixed pattern noise as well as its correction. This is particularly relevant for logarithmic sensors, which suffer a stronger impact from uncorrected fixed pattern noise than traditional linear sensors. However, results from chapter 4 showed that the absolute contrast error for a high dynamic range monochrome image from a logarithmic sensor would be at least 5% over 4 decades while about 5 decades would be covered at below 2% if the relative contrast error were considered. It is, however, unclear how these figures translate to our understanding of colour and specifically in the processing and reproduction of colour images.

As a result, one of the questions addressed in this chapter is how to relate the absolute and relative contrast errors discussed in earlier chapters to the ΔE_{76} error measure. This will emphasise thinking of image quality after fixed pattern noise correction in terms of a perceptual model with three degrees of freedom, rather than as a one-dimensional attribute. Related to this is the effect of the fixed pattern noise correction method on colour quality, and whether the model error arising from the two parameter model extraction technique is significant to acquiring the colour image standard of less than 3 ΔE_{76} units adopted in section 5.3.1. This is important because it shows limits of the correction procedure and the

effects of quantisation on colour quality.

The effect of contrast error on the quality of colour images is first investigated by evaluating the visual effect of different changes on RGB values of the Munsell colours of section 6.2. The distinction between absolute and relative contrast in relation to colour is also discussed. With encouraging results and sufficient background to the way colour information is manipulated, a MATLAB simulation is executed on $1,232$ displayable Munsell colours through a logarithmic sensor - section 6.3.1 - to simulate the effects of model error on colour quality. The effect of the two parameter extraction routine on the position of the colour within a standard CIE colour gamut is investigated over 10 decades of illumination at specific points of interest - section 6.3.2. For a better understanding of the visual effect on colour perception, these results are then shown as changes in the perceptually uniform CIELAB colour space. Both ΔE_{76} and ΔE_{00} units are shown for comparison. Following the low perceptual errors in the mid regions of the illumination range, analysis explaining the validity of this result is then detailed in sub-section 6.3.3.

However, to model the effects of the imaging processes on colour accurately, the other sources of noise have to be added to the simulation. In section 6.4, the results of modelling fixed pattern noise are presented against CIELAB units. Similarly, the ΔE_{76} and ΔE_{00} units are shown to provide a comparison of the results. In section 6.4.2, ADC quantisation noise is also included in the discussion. As in the case with monochrome images, the effect of the analogue to digital converter on the user's ability to perceive colour quality is evaluated at different bit sizes in addition to residual FPN.

6.2 Colour and Contrast

So far, image quality has been discussed in the context of two separate scenarios. When considering monochrome images, the just noticeable difference of the human visual system was initially used to describe an absolute contrast error when logarithmic sensor output was used to estimate a photocurrent from the pixel imaging a scene. This was modified to a relative contrast error in accordance with the adaptive properties of the HVS, which leads to insensitivity to absolute illumination. However, once colour images are captured it is necessary to consider the accuracy with which these colours are captured. The ΔE_{76} error measure which takes into account the visual spectral response to the three primaries was then adopted in the last chapter as a measure of accuracy.

As a result when fixed pattern noise was corrected in monochrome images, the results of employing the two parameter extraction procedure yielded at least 5% absolute contrast error or below 2% relative contrast error. This means that for colour images, it is possible to say that each of the red, green or

blue primaries imaged through the same correction scheme could have up to 5% absolute or 2% relative contrast error respectively. It is difficult to tell what amount of colour error will result from this amount of residual model error in the individual primaries. The importance of being able to quantify the relationship between contrast error and CIELAB error is that we can reliably tell the expected colour quality without including all the other processes that limit colour fidelity even further.

Before an investigation into this relationship can commence, it is vital to understand the interplay between relative and absolute contrast with regard to colour. While the relative contrast error meant that, in monochrome images, it did not matter if an estimated photocurrent at a pixel location was incorrect as long as the estimated photocurrents at other locations were systematically incorrect, the same might not be entirely true for colour images. When the estimated red, blue and red components of a sample colour all change, or are incorrect relative to each other, by a constant multiplicative factor, the chroma component of the sample colour will stay the same since the ratios between the bands will be unchanged - that is the xy component (chroma) of the xyY space is constant and the dominant hue is left unaltered. However, the lightness component and therefore the overall perception of colour will have been distorted. Therefore, at best relative contrast preserves chroma, but overall it has little significance in maintaining colour fidelity of an imaging process.

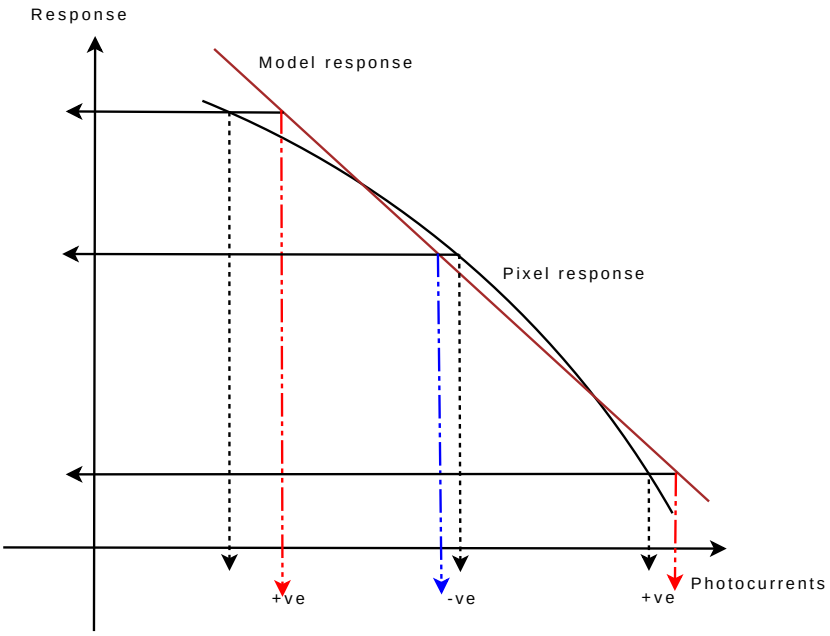

Figure 6.1: A simple illustration of model error and the sign or polarity of the possible estimated photocurrents and absolute contrast error.

To study the effect of absolute contrast in each of the primaries of a colour's perception, a simulation

was performed where errors were progressively introduced to the red, green and blue bands of the Munsell colours and the CIELAB error units calculated. First, all the non-displayable colours of the Munsell set, that is those whose red bands are negative, were removed since CIELAB is a perceptual space. In deciding the amount and polarity of the error to be inserted to the bands, experience in the modelling of logarithmic pixels and figure 2.8 show that absolute contrast errors of about 5% can be reasonably expected. In addition, the estimated red, green or blue contrast errors usually fall in the same region of the absolute error curve and are commonly of the same sign - a fact attributed to the low ratios between the colour components as illustrated in figure 5.6.

Figure 6.1 shows the two parameter model superimposed over an exaggerated logarithmic response curve with the estimated photocurrents represented as the blue and red dotted lines intercepting the model response while the actual photocurrents are shown as black ones. It is clear that these estimated currents usually create either all positive or all negative absolute contrast errors depending on the region in question. This follows from the low ratios between the colour bands exhibited by the sample colours in section 5.4.1. The RGB band errors therefore tend to follow in the same negative direction, are low and close together in the linear region but have large positive errors in the dark and moderate inversion regions.

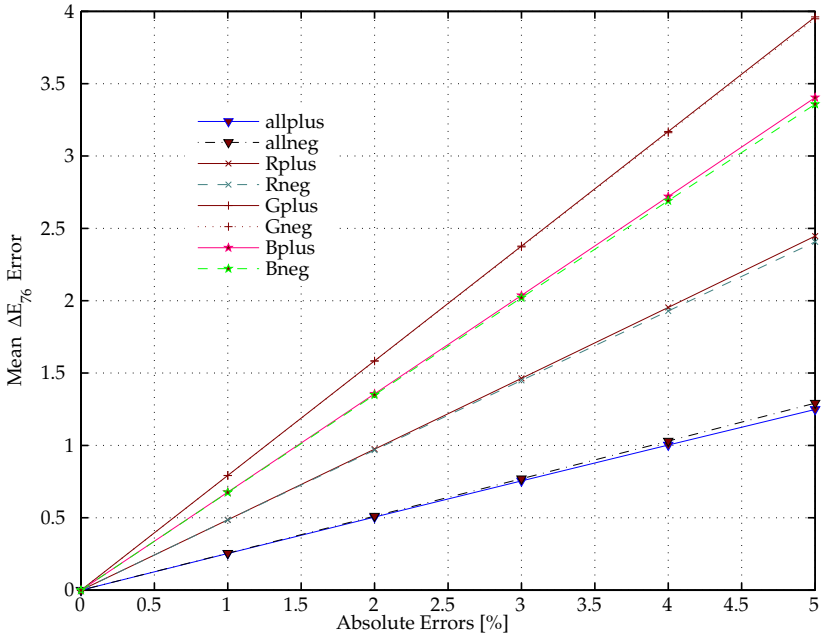

Figure 6.2: The effect of different combinations of the sign of the contrast error, created by introducing errors in RGB, on the mean ΔE_{76} CIELAB error.

As a result, the simulation involved either subtracting or adding the same 1%, 2%, 3%, 4% and 5%

errors in the red, blue and green bands of all the displayable Munsell colours. This was done one band at a time and then on all the three to show the different combinations of positive and negative errors and thus the worst combination or highest CIELAB error. Figure 6.2 shows the different permutations against the mean ΔE_{76} error of all Munsell colours where a positive absolute contrast is designated by a "plus" and the converse represented by "neg" suffixed to the colour band.

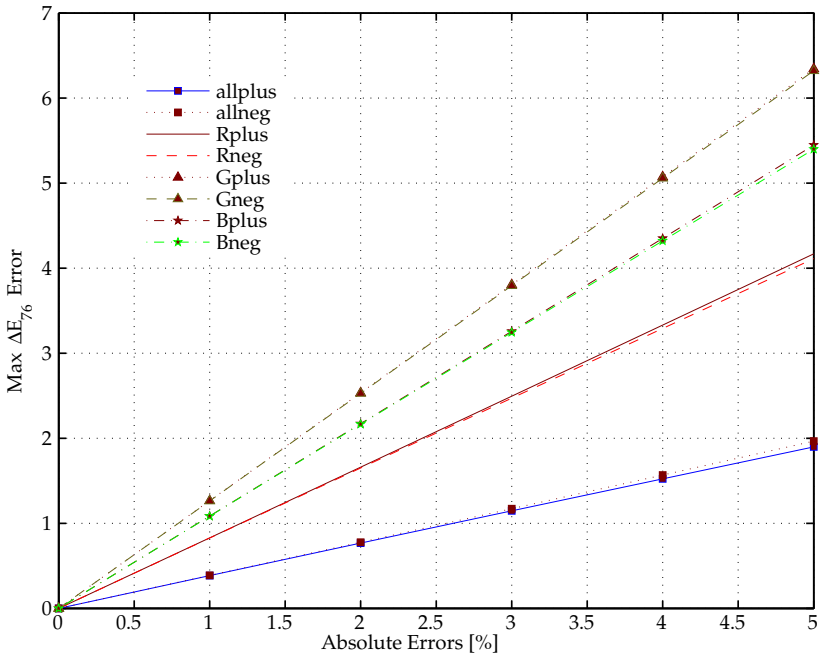

Figure 6.3: The effect of different combinations of the sign of the contrast error, created by introducing errors in RGB, on the maximum ΔE_{76} CIELAB error.

The figure reveals that the polarity of the error is insignificant with respect to colour error unlike the magnitude of the error itself. This follows from equation (5.10). In fact, as the contrast error increases in the RGB bands, the perceived colour quality drops proportionally. This is thought to arise from the fact that for very small changes, the definition of CIELAB dictates a linear perceptual relationship, equations (5.7) - (5.9). Inspite of this, the worst case CIELAB error occurs when the green channel is changed while the rest are kept constant; possibly, because the eye is most sensitive to green changes. Interestingly, the lowest colour error is achieved when all the colour bands have error of the same sign introduced in them. This is because, the human visual system, on which the Munsell colour convention is based, is an opponent colour system hence colour is registered as differences between red and green, blue and and yellow and finally black and white. The differences between errors in all the bands then lead to lower perceptual ΔE_{76} units. It is therefore reasonable that since all the bands are shifted during FPN correction and that for contrast errors of up to 5% in RGB, average perceptual errors of about $1-2$ units

can be expected. The same logic is applicable in the results of figure 6.3, except for the higher values since the maximum error is registered in this case. Even then, just about 2 ΔE_{76} units are obtained as the maximum error when an absolute contrast of 5% is present in the bands.

Generally, it can be expected that the absolute errors of up to 5% are tolerable in the photocurrents estimated in the purely logarithmic region of the illumination range, leading to about $1-2$ ΔE_{76} units. This level of accuracy is similar to the accuracy that is obtained with typical FPN correction. Therefore, these simulations offer a good indicator of probable colour accuracy and provide a basis to undertake a detailed model of the effect of model error on colour accuracy.

6.3 Colour and Model Error

6.3.1 Matlab Simulation

To fully investigate the impact of the contrast error on colour quality over a wide dynamic range, a MATLAB simulation using the Munsell colours is performed over 10 decades of illumination. With the intention of mimicking the parameter extraction process using the two parameter model, a function was developed that accepts RGB Munsell colour values that are to be imaged. Thus, by comparing the CIELAB values of the Munsell colours before and after parameter extraction at all illuminations, it is hoped that the ΔE_{76} colour errors will show whether the target colour accuracy has been achieved.

In this simulation, only $1,232$ colours of the Munsell data set were used since it is not logical to include non-displayable colours in a calculation involving CIELAB, which is a perceptual measure. However, the first problem is that CIELAB requires a reference white colour to determine colour error. None of the colours in the Munsell set is white. Secondly, the lightness component of colour is only a low dynamic range representation of the amount of luminance a particular colour is inherently associated with (section 5.4.1). Therefore the transformation of the RGB colour values over a high dynamic range would require the linearisation of the colours on a $[0-1]$ scale leading to colours at the mid to lower light levels being misrepresented.

To combat these problems, an equivalent $D65/2$ white colour was added to the Munsell colour set. The choice of $D65/2$ lighting is solely because it most closely matches daylight, an environment in which sensors are usually operating in. The addition of a white colour also meant that by viewing all colours at every illumination, it was possible to reference all colours to the brightest possible colour at that illumination and therefore have a $[0-1]$ range, required for making colour space transformations.

In other words, to represent all the colours in a HDR context, it meant simulating a logarithmic pixel as being illuminated by all the colours relative to the brightest white colour for each photo-illumination.

As discussed in earlier chapters, two carefully chosen photocurrents and responses at those currents are required in the parameter extraction procedure. It is these parameters that are used to acquire the estimates of the initial photocurrents. First, to create a high dynamic range response, a logarithmic pixel was simulated from 1×10^{-16}Amps to 1×10^{-6}Amps in 0.1 decade intervals. This was possible using typical characteristic values of the offset a_j, gain b_j, leakage current c_j and inversion current d_j parameters discussed in chapter 4. Furthermore, to create the closest possible match to a typical logarithmic pixel response, Gaussian noise with statistical properties equal to the difference between CADENCE simulations[1] and the response of the four parameter model, was added. With the knowledge of the photocurrents and the c_j and d_j parameters, the offset and gain parameters of the pixel were then extracted using two well-chosen data points - 1×10^{-13}Amps and 1×10^{-9}Amps.

Now, with a total of $1,233$ colours, the pixel was assumed to capture one colour at a time at each illumination for all the colours. Thus a red, green and blue photocurrent was generated for each colour from 1×10^{-16}Amps to 1×10^{-6}Amps. Using the offset and gain parameters, extracted earlier, the estimated red, green and blue photocurrents from using the two parameter model response were determined. These RGB estimates of the photocurrents for all colours after parameter extraction, were then normalised using the equivalent red, green and blue photocurrents of the white illuminant so as to produce a standard $[0-1]$ low dynamic representation. This allowed for standardisation but mostly permitted transformation to the xyY and CIELAB colour spaces for comparison.

6.3.2 Simulation Results

By comparing the xyY values of the colours before and after applying the two parameter extraction routine on the typical high dynamic range response, it is possible to know how colour accuracy is affected by model error. This is possible when the chromaticity coefficients are calculated from the respective XYZ values and the subsequent positional shifts are observed on a chromaticity diagram.

In figure 6.4, a random sample of 50 Munsell colours are shown on a chromaticity diagram identified at 1×10^{-14}Amps which corresponds to a decade before the dark current of the logarithmic pixel. Each line represents a shift in chroma where the initial small dot is the previous colour position while the "cross" is the position after applying the two-parameter model on the Munsell colour in question.

[1]The difference between the four parameter model and typical simulated responses per illumination was found to be approximately gaussian with a near zero mean but a standard deviation of 1.5×10^{-5}.

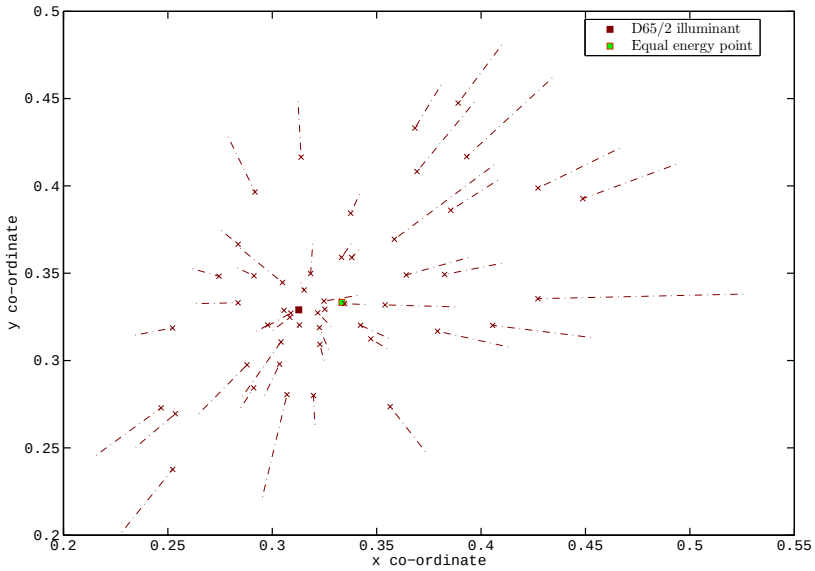

Figure 6.4: The effect of the simple model parameter extraction on the hue of a sample of 50 Munsell colours, one decade away from the leakage current photo illumination (1×10^{-14}Amps), shown as part of a chromaticity diagram.

It is clear that at this illumination, most of the colours have a severely distorted chroma hence the perception of the original colours will be substantially different. It can also be comfortably deduced that the ratios between the colour bands that determine the dominant hue have been altered by imaging the colours through the two parameter model for this illumination. These hue shifts can be explained by the inaccuracy of the two-parameter model in a region where leakage currents are significant. Since the two parameter model is inadequate for modelling logarithmic behaviour at this illumination, the estimates for the red, green and blue colour components are equally misrepresented and hence the poor rendition.

In spite of the varying chroma shifts, all the fifty colours tend to the $D65/2$ white illuminant which is represented on the same figure. The direction of the shift is dictated by the fact that at this dark photo-illumination all the red, green and blue photocurrent estimates are registered closer to each other or as more similar values - figure 6.1. This is close to a gray or white colour. However, as the chromaticity diagram is independent of lightness, the shift towards the white colour at the dark current, actually represents a reduction in the saturation of the colour and a drift towards the darkest shades of gray on the lightness plane orthogonal to the chromaticity diagram. In the extreme case and because of the very low sensitivity of the pixel at very low illumination, the pixel output is constant for all colours leading to the same output which is black but will appear as the $D65/2$ illuminant point in the lightness independent

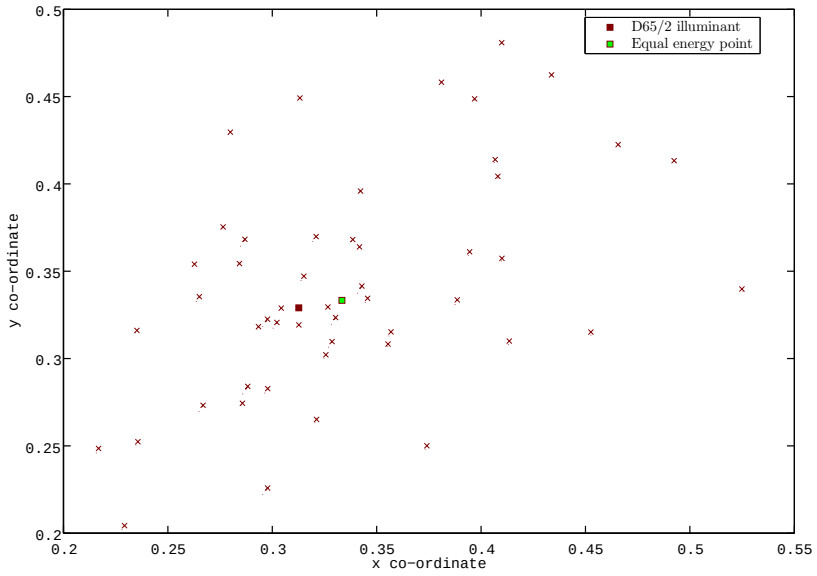

Figure 6.5: The effect of simple model parameter extraction on the hue of a sample of 50 Munsell colours at an illumination in the purely logarithmic region (1×10^{-11}Amps), shown as part of the chromaticity diagram.

chromaticity diagram.

When the positions of the same fifty colours are evaluated at an illumination in the middle of the purely logarithmic region, 1×10^{-11}Amps, the shifts in chroma appear almost insignificant as shown in figure 6.5. Understandably, the two-parameter model accurately predicts the typical logarithmic pixel response thereby correctly estimating the red, green and blue photocurrents. This inevitably leads to very low absolute and chroma errors. The residual model error seems insufficient to cause significant hue differences in the sampled Munsell colours.

An investigation of how the colours are affected at high photocurrents reveals a different result. At a decade before moderate inversion, 1×10^{-8}Amps, the two parameter response is invalid for predicting accurate pixel response. The resulting high absolute errors produce random colour shifts as shown in figure 6.6. The pattern appears to be that the colours increase in saturation, as evidenced by their shift away from $D65/2$ white. This is wholly attributed to the use of a two parameter logarithmic model in a region where, as discussed in chapter 3, a square law relationship begins to take shape. Consequently, the steep slope in this region leads to the red, blue and green colour components being altered differently for different colours.

Although all the figures in this section show differences in the chroma that occur after the application

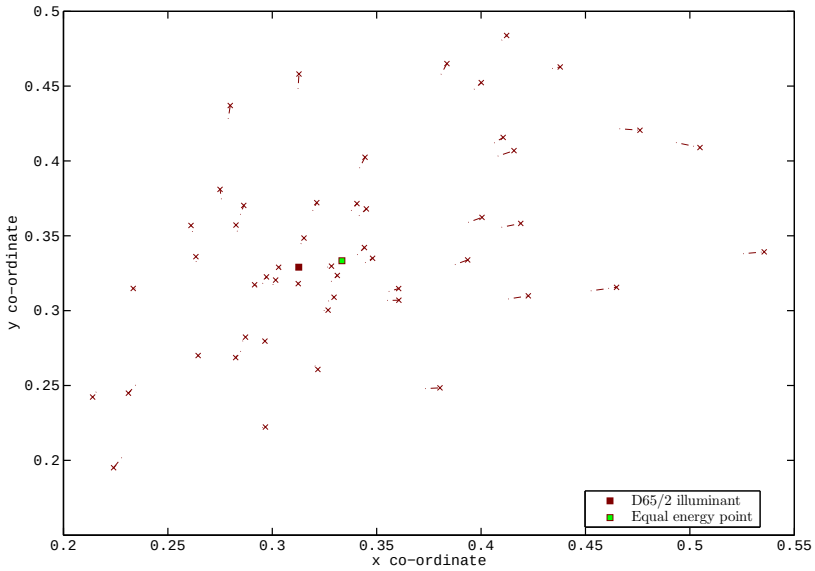

Figure 6.6: The effect of simple model parameter extraction on the hue of a sample of 50 Munsell colours at an illumination one decade from moderate inversion (1×10^{-8}Amps), shown as part of the chromaticity diagram.

of the two parameter model, a measure of how perceptually different the colours appear is not represented. This is because the chromaticity diagram is a perceptually non-uniform space [3]. Consequently, it is not possible to determine whether the shifts that appear as long lines are actually significant perceptual colour shifts. Subsequently, the effect of absolute errors has to be evaluated in CIELAB if the true effect of residual model error are to be determined.

Figure 6.7 shows the distribution of ΔE_{76} values of all the colours over the entire illumination range. While the vertical axis shows the Munsell indices of only the displayable colours, the horizontal axis is representative of the photocurrent a white reference illuminant would generate in a pixel. The figure also confirms previous findings in figures 6.4 - 6.6 that suggested that colour errors at the lower illumination range are higher than those in the middle range and but not as bad as the high current end. In fact, the very high errors at the lower illumination are witness to the loss of colour accuracy, with CIELAB values often greater than 50 units and not confined to particular hues as evidenced by the even spread along the Munsell index.

However, from this figure it is difficult to tell the precise levels of colour accuracy for the vast majority of photocurrents especially in the mid current region. This is because most of these values broadly fall under 20 ΔE_{76} units, and illuminations under which smaller colour errors may exist are not readily visible.

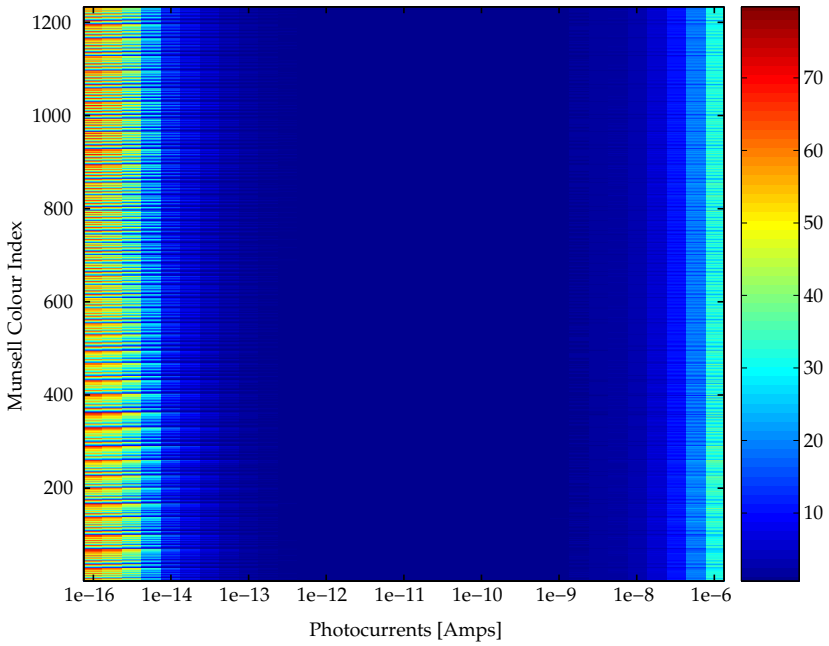

Figure 6.7: A representation of the ΔE_{76} CIELAB error range and distribution after two parameter extraction over the entire 10 decade illumination range.

Figure 6.8 shows the cumulative distribution functions of the ΔE_{76} errors at specific photocurrents along the illumination range. From about 1×10^{-14} Amps until 1×10^{-8} Amps, it is possible to see that at least 50% of the Munsell colours yield under 10 ΔE_{76} units, with the least colour error experienced in the 3 decade region from 1×10^{-12} to 100×10^{-12} Amps. Here, nearly all the colours have less than 2 ΔE_{76} units.

Nevertheless, to find out the precise dynamic range over which acceptable CIELAB errors may result from the application of the two parameter model, the mean colour error values per photocurrent were determined and presented in figure 6.9. This figure shows both ΔE_{76} and ΔE_{00} errors that represent the worst and best case metrics respectively. In both cases, the colour error profile follows that from figure 6.7 and for approximately 5 decades of illumination, model error accounts for less than 3 error units. This is considered hardly perceptible colour quality - table 5.1 and section 5.4.2. It is also sufficient for average natural scenes and a variety of commercial or industrial scenes but ultimately implies that the residual absolute contrast values when employing the two-parameter model during logarithmic imaging, are indeed manageable.

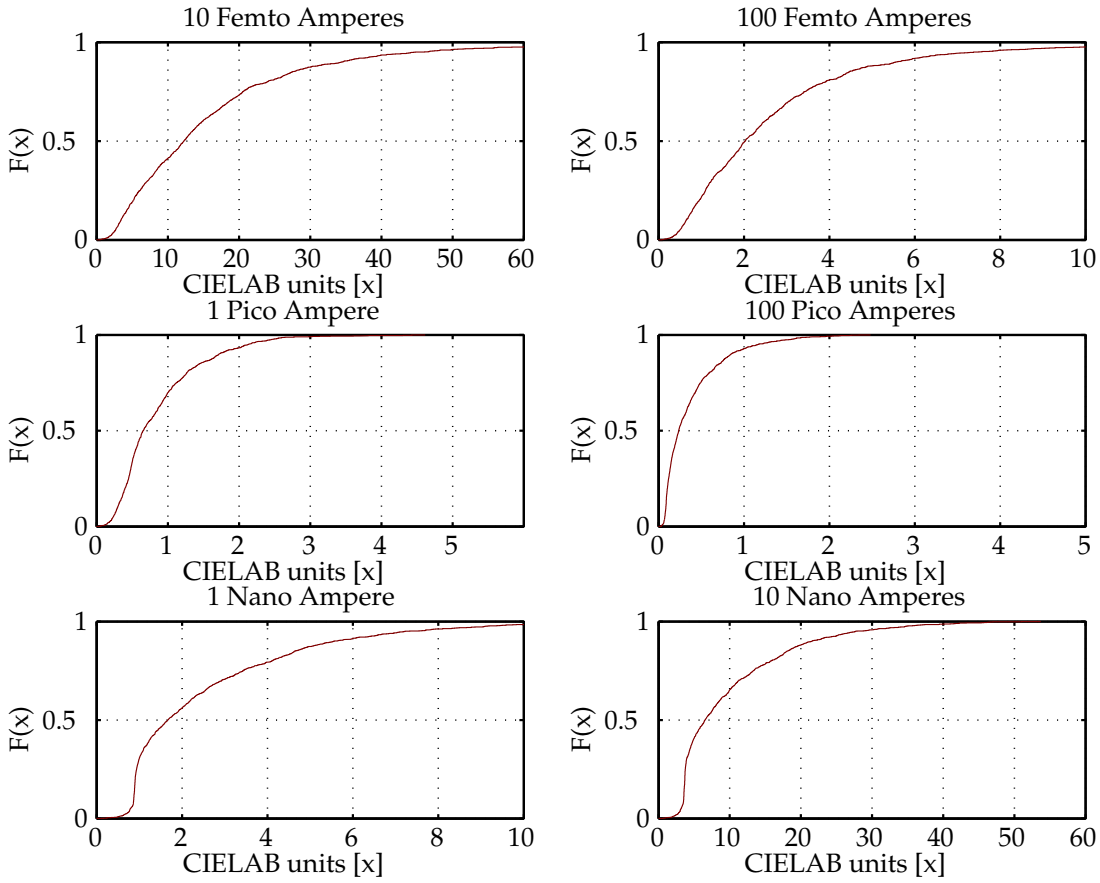

Figure 6.8: The cumulative distribution function of the ΔE_{76} errors at different photocurrents along the illumination range.

6.3.3 Error Analysis

The very low ΔE_{76} error units obtained during the investigation of colour accuracy in section 6.2, fall within the predictions of the earlier simulations in section 6.2. They may be further understood by taking an analytical approach, based on the definition of the CIELAB colour space components. It is possible to express the model error in any of the estimated red, green and blue photocurrents in terms of the absolute contrast α

$$\begin{aligned} \alpha &= \frac{\Delta I}{I} = \frac{I_{act} - I_{est}}{I_{act}} \\ \therefore \quad I_{est} &= I_{act}(1-\alpha) \end{aligned} \tag{6.1}$$

where I_{est} is the estimate of the actual photocurrent I_{act}, after applying the two parameter model to the pixel response for each primary. After transforming the RGB photocurrents to XYZ and then to the

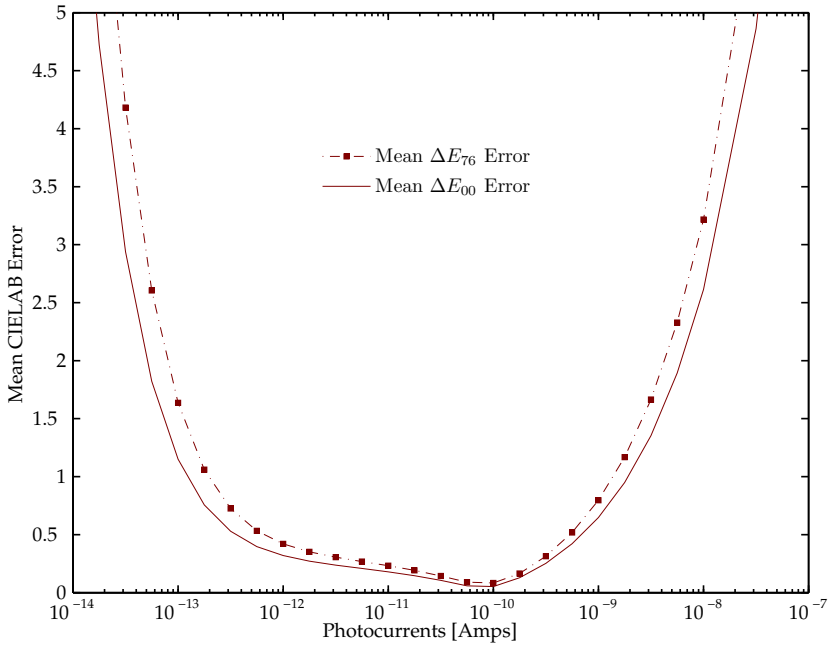

Figure 6.9: The mean ΔE_{76} and ΔE_{00} CIELAB error values for perceptual accuracy over $1,232$ displayable Munsell colours. The ΔE_{76} measure is considered the worst case while the latter is the best measurable option.

CIELAB colour space, the lightness component L before and after parameter extraction, can be written in terms of absolute contrast by applying equation (6.1) to get equation (6.2); where the subscripts $_K$ and $_W$ designate the Munsell colour and the $D65/2$ illuminant respectively. Note that only the first part of equations (5.7) - (5.9) that is $f(t) > 0.8856\%$, has been taken into consideration from the CIELAB equations (5.7) - (5.9) , since most colours in averagely illuminated scenes fulfill this condition.

$$\begin{aligned} L_{act} &= 116\left[\frac{Y_{Kact}}{Y_{Wact}}\right]^{\frac{1}{3}} - 16 \\ \therefore \quad L_{est} &= 116\left[\frac{Y_{Kest}}{Y_{West}}\right]^{\frac{1}{3}} - 16 \\ &= 116\left[\left(\frac{1-\alpha_K}{1-\alpha_W}\right)\left(\frac{Y_{Kact}}{Y_{Wact}}\right)\right]^{\frac{1}{3}} - 16 \end{aligned} \tag{6.2}$$

α_K represents the RGB absolute contrast error transformed in the XYZ space while α_W represents the same measure of the reference white colour. Similarly, the redness-greenness a plus the yellowness-

blueness b components of the corresponding colour can be written as;

$$\begin{aligned} a_{act} &= 500\left[\left(\frac{X_{Kact}}{X_{Wact}}\right)^{\frac{1}{3}} - \left(\frac{Y_{Kact}}{Y_{Wact}}\right)^{\frac{1}{3}}\right] \\ \therefore \quad a_{est} &= 500\left[\left(\frac{X_{Kest}}{X_{West}}\right)^{\frac{1}{3}} - \left(\frac{Y_{Kest}}{Y_{West}}\right)^{\frac{1}{3}}\right] \\ &= 500\left[\left(\frac{(1-\alpha_K)X_{Kact}}{(1-\alpha_W)X_{Wact}}\right)^{\frac{1}{3}} - \left(\frac{(1-\alpha_K)Y_{Kact}}{(1-\alpha_W)Y_{Wact}}\right)^{\frac{1}{3}}\right] \end{aligned} \tag{6.3}$$

and

$$\begin{aligned} b_{est} &= 200\left[\left(\frac{Y_{Kest}}{Y_{West}}\right)^{\frac{1}{3}} - \left(\frac{Z_{Kest}}{Z_{West}}\right)^{\frac{1}{3}}\right] \\ &= 200\left[\left(\frac{(1-\alpha_K)Y_{Kact}}{(1-\alpha_W)Y_{Wact}}\right)^{\frac{1}{3}} - \left(\frac{(1-\alpha_K)Z_{Kact}}{(1-\alpha_W)Z_{Wact}}\right)^{\frac{1}{3}}\right] \end{aligned} \tag{6.4}$$

Now, the residual absolute contrast errors as seen in chapter 2 range from very low values in the purely logarithmic region to very high values in the leakage and moderate inversion regions of illumination. In the mid current range, contrast error is in the range $0-5\%$. Here, an approximation of the binomial series

$$(1+x)^n = 1+\frac{nx}{1}+\frac{n(n-1)x^2}{2!}+\frac{n(n-1)(n-2)x^3}{3!}+\cdots$$

can be applied to simplify the expressions for the estimated CIELAB components. Since $\alpha_K << 1$ and $\alpha_W << 1$ the higher powers of the bracketed terms involving the contrast error, in equations (6.2), (6.3) and (6.4) can be ignored to get

$$\frac{(1-\alpha_K)^{\frac{1}{3}}}{(1-\alpha_W)^{\frac{1}{3}}} \approx \frac{1-(\alpha_K/3)}{1-(\alpha_W/3)} \tag{6.5}$$

Thus, despite the different but low varying values of α_K and α_W from the separate colour bands and white respectively, equation (6.5) tends towards unity. In other words $L_{est} \longrightarrow L_{act}$, $a_{est} \longrightarrow a_{act}$ and $b_{est} \longrightarrow b_{act}$. As a result the calculation of colour difference using equation (5.10), then reduces to very low ΔE_{76} values since $\Delta L \longrightarrow 0$, $\Delta a \longrightarrow 0$ and $\Delta b \longrightarrow 0$ in this region. The estimated photocurrents have been correctly predicted and the colour ratios between the bands kept almost constant. Conversely, outside the purely logarithmic region, the contrast values are high since the model inaccurately predicts the estimated photocurrent hence the simple binomial series approximation cannot be used. The distorted lightness and chromatic components of CIELAB vary significantly from the original values thereby

leading to high ΔE_{76} colour error.

6.4 Colour and Noise Simulation

The simulations above have shown that by using the two parameter extraction technique in a logarithmic colour sensor, the residual model error leads to negligible colour error in the purely logarithmic region of illumination. However, the simulation needs to be extended to cover the potentially degrading effects of noise. Various sources of noise cannot be all accurately simulated, but the addition of the most prevalent can highlight the most likely outcome when characterising real logarithmic colour sensors.

6.4.1 FPN Effects

The dominant form of noise that adversely affects logarithmic sensor performance is spatially varying fixed pattern noise as extensively discussed in previous chapters. It is therefore imperative that FPN is added to the analysis of colour error if conclusions that are more meaningful are to be drawn. Therefore, the simulations detailed in section 6.3.1 were extended to include FPN noise using typical characteristic values of the variation of offset, gain, leakage and moderate inversion parameters, obtained from a real logarithmic sensor - table 4.1. The mean value of each of these parameters (determined by non-linear optimisation) as well as their variation was replicated by randomly adding a Gaussian noise distribution to the responses of $1,000$ pixels in an effort to mimic fixed pattern noise.

The same 1232 Munsell colours as well as a reference white colour were applied to each of the different pixels, one colour at a time over the entire 10 decade illumination range. Again the offset and gain parameters were first extracted, this time for each of the different $1,000$ pixel responses at the same two calibration points. These extracted offset and gain parameters were then used to obtain the estimated red, green and blue photocurrents for each colour at each pixel and at each illumination. After normalising the estimates of the primaries against the white reference colour in the set, the CIELAB values before and after parameter extraction and FPN correction were obtained.

Figure 6.10 shows the result of the simulation, with the curves revealing the mean ΔE_{76} and ΔE_{00} errors over all pixels and overall $1,233$ colours compared with those without any fixed pattern noise effects. The ΔE_{00} error curves have lower error than the ΔE_{76} ones but without a significant difference. In all cases however, more accurate modelling in the purely logarithmic region still creates lower colour error than outside the mid-current region. It can be noted that for both error metrics the results of introducing FPN, and subsequently correcting for it, gives results that are nearly the same as when no fixed

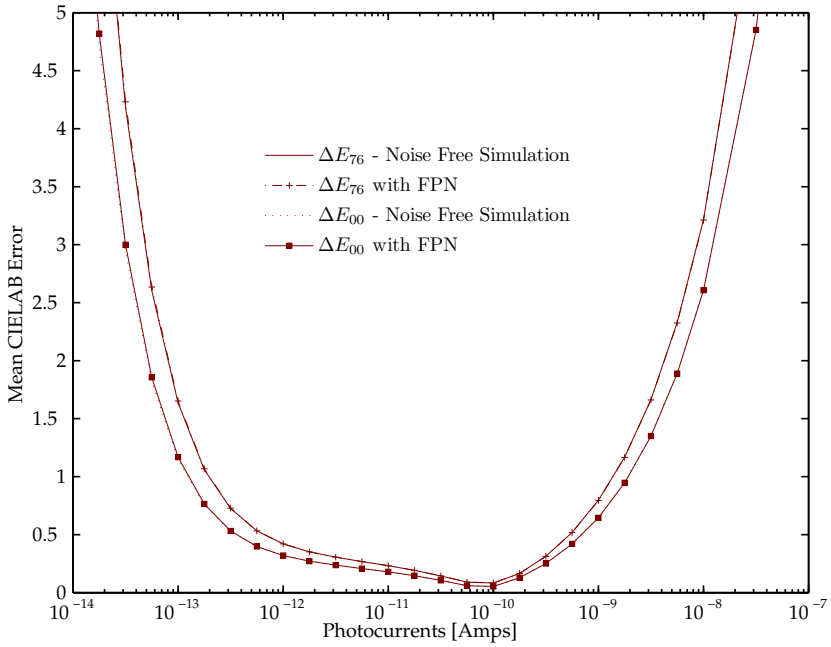

Figure 6.10: Results of a high dynamic range simulation of the colour performance of a logarithmic imager after modelling fixed pattern noise correction.

pattern noise effects was applied. The very close results in the purely logarithmic region are no surprise; but outside this region the CIELAB error arising from the poorly estimated photocurrents is known to overshadow the colour error caused by the uncorrected variations in the leakage and inversion current parameters, c_j and d_j respectively. This means despite adding typical variations in all the four parameters to mimic FPN, the correction procedure accurately removes mismatch leaving only the model error that was previously simulated. The inference is that the offset and gain variations are all but eliminated while the variations in the c_j and d_j parameter lead to insignificant colour error.

6.4.2 Quantisation Effects

The digitising process of the analogue to digital converter always introduces quantisation noise that affects the precision of the signal output of the pixel. The amount of noise associated with this process in any sensor depends on the chosen bit length and output signal swing. For typical logarithmic pixels, low light responsivities of about 300mV total voltage swing make the number of bits at the ADC, the prime determining factor for the amount of quantisation error likely to be encountered by the sensor.

Thus, in addition to model and residual fixed pattern noise, the effects of the quantization process at analogue to digital conversion have to be considered in the MATLAB simulation in order to gain a more complete understanding of the effects of the two parameter model on colour accuracy. As before the

responses of the 1,000 pixels were modelled to include the same typical fixed pattern noise variations and at the same levels as in previous simulations. This time however, the pixel responses were rounded off to 0.2mV, 0.4mV, 0.8mV, and 1.6mV during each simulation. These values reflect the step sizes that are required when a typical 50mV/decade logarithmic pixel sensor, experiencing 100mV peak to peak fixed pattern noise is employed with either an 11, 10, 9 or 8 bit ADC respectively over a six decade illumination range. Two-parameter fixed pattern noise correction was then carried out on the rounded values to obtain the estimated red, green and blue primaries before determining the ΔE_{76} CIELAB values for the various quantisation levels.

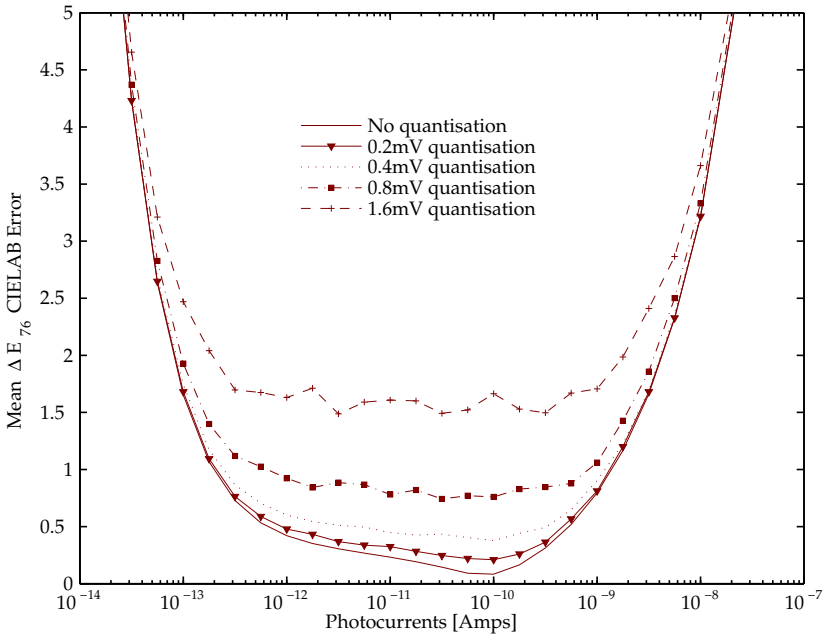

Figure 6.11: An investigation of ADC quantisation effects, on the simulation of logarithmic imaging of the Munsell colours. This simulation includes fixed pattern noise effects as well.

In figure 6.11, the increase in analogue to digital converter step size versus the detrimental effects of colour perception after fixed pattern noise correction is highlighted. The figure shows that with a doubling of the level of quantisation, there is a corresponding proportional increase of the ΔE_{76} error units in the purely logarithmic region. A doubling of step size from 0.4mV to 0.8mV causes an approximate 0.5 increase in ΔE_{76} while a change in bit resolution from 0.8mV to 1.6mV leads to about a ΔE_{76} unit change. This can be explained using the link between contrast and CIELAB illustrated in figure 6.2.

The 0.8mV step change in quantisation level equates to a maximum rounding noise margin of half that step size. But because from quantisation theory the root mean square value is a factor of $\sqrt{12}$ of the step level, it is equivalent to only about 0.23mV rms or approximately 1.1% absolute contrast for a

50mV/decade logarithmic pixel sensor. From figure 6.2, it is possible to see that the corresponding ΔE_{76} error could be up to a unit error - the same change encountered in figure 6.11.

Coupled with the increase in colour error in figure 6.11, there is also a slight decrease in dynamic range over which typical logarithmic behaviour can be claimed, for a specified colour accuracy. For a colour difference error of 1 unit and with a 0.4mV step size, 4 decades of logarithmic operation can be claimed whereas at twice this bit size, only about 3 decades are achievable at this accuracy. In table 6.1, the relationship between the quantisation level and the required bit length to cover a 4 and 5 decade illumination range is highlighted. Note that the bit sizes are still based on a 50mV/decade logarithmic pixel with and a 100mV fixed pattern noise range. Overall, logarithmic sensors can achieve 5 decades of illumination under 3 ΔE_{76} units and even an extra decade if colour error under 5 ΔE_{76} units is tolerable.

Dynamic Range	Quantisation Level	0.2mV (11 bits)	0.4mV (10 bits)	0.8mV (9 bits)	1.6mV (8 bits)
4 decades	ΔE_{76} value	1.5	1.5	1.5	2.0
5 decades	ΔE_{76} value	2.5	2.5	2.5	3.0

Table 6.1: A comparison of the ADC bit length for particular image colour quality levels and associated degrees of quantisation.

The fairly low number of bits to accurately image a 5 decade dynamic range scene also emphasises the logarithmic sensor's main advantage in HDR imaging. For example, to achieve average frame rates of 30fps, an SVGA (800×600 resolution) sensor having a 0.4mV step size ADC will need a minimum data rate of 144MHz. On the contrary, linear sensors attempting to achieve similar accuracy over five decades without multiple sampling, will definitely require more than 10 bits and a higher data rate[2]. This is even without the extra cost of circuit considerations needed for HDR imaging and achieving the level of colour accuracy.

6.5 Summary

This chapter has extended the results of correcting fixed pattern noise in monochrome images to high dynamic range colour scenes. Because these images compose the vast majority of user images, this study was deemed significant if logarithmic sensors are to be useful in many high dynamic range-imaging applications.

Using basic but essential knowledge of colour reproduction from chapter 5, the link between contrast error and colour quality was first investigated. By studying the effect of absolute error in each of the

[2]Linear sensors are generally known to require over 10 bits to image about 3 decades of illumination at about 10% accuracy.

constituent colour bands, the limitations of model error on the overall perception of colour were deemed tolerable. This led to a MATLAB simulation using the standard set of Munsell colours over a 10 decade range of photo-illumination.

A white reference colour was first added to the Munsell colour set as a way of transforming RGB to the perceptual CIELAB, but also as a means of normalising all the colours at any illumination. Using typical values of the a_j, b_j, c_j and d_j parameters, the response of a pixel over the illumination range allowed for the extraction of offset and gain from the two parameter model. Modelling for colour assumed all the colours imaged by a pixel one at time at each illumination. Thus, using the extracted offset and gain parameters, the estimated red, green and blue photocurrents were obtained for each illumination. After referencing all the colours to the reference white at that illumination, the estimated RGB colours were converted to CIELAB before comparing to their original RGB values.

By investigating the shifts in the positions of the colours on a chromaticity diagram at different illuminations, it was possible to begin to determine the effects of absolute error on colour. Less chroma shifts were observed in the purely logarithmic region while worse ones were present outside this region. CIELAB values confirmed this trend with colour error falling below three ΔE_{76} or ΔE_{00} units for most of the photocurrent range. An analytical investigation into the low CIELAB values then confirmed earlier predictions in section 6.2. Thus, residual absolute errors are actually manageable since they do not adversely affect colour perception.

However, to mimic real sensor noise conditions, typical fixed pattern noise variations were added to the simulation over $1,000$ pixels. Results revealed that the addition of FPN and subsequent image correction did not change the perception of colours over the entire range. Then, ADC quantisation noise effects were also added to the MATLAB simulation to match the typical imaging process even closer. Digitisation noise was therefore added to the previous simulations for the same $1,000$ pixels. As the step size was increased, a corresponding drop in colour accuracy was realised. Inspite of this, a 5 decade range was still obtainable at less than 3 ΔE_{76} units of colour accuracy for a 10 bit ADC.

Overall, this work reveals the very high potential for high dynamic range monochrome and colour images from a logarithmic sensor that performs simple fixed pattern noise correction. The resulting images are perceptually indistinguishable over at least 5 decades making them ideal for most contemporary high dynamic range image applications.

Chapter 7

HDR Display

7.1 Introduction

The preceding chapters have shown that high dynamic range images captured with a logarithmic sensor can be corrected for fixed pattern noise, achieving both uniformity and perceptibility of colours at levels comparable to that of the human eye. This is true for approximately 5 decades of illumination change of the imaged scene for both monochrome and colour images. In these high dynamic range scenes, the preservation of feature and colour quality after image capture and processing is equally as vital. This is because for the majority of applications, the display of these images is used for decision-making and / or rendition purposes. Processes such as quality control in robotics, machine and image manipulation for the automobile industry and recognition routines for security applications, all rely on the best reproduction of captured scenes on display screens. Thus, the image sensor can no longer be considered in isolation but as part of an image handling chain, which has a role to play in subsequent image display.

The delivery of images has mostly taken the form of screen display and print media although the former is by far a more common output method. However, the dynamic range of both visual and print media is limited to a few decades, typically just above 2 decades because of the inherent device physics of display technologies. Consequently, displaying or printing wide dynamic range images, sometimes extending over 6 decades, commonly results in over saturation of the very bright parts of the scene and / or the under exposure of the darker parts of the same scene. This loss of detail in some parts of an image is not desirable for any application that seeks high image fidelity.

In this chapter, the theoretical considerations that are required to display high dynamic range log-

arithmic images on standard 8 bits screens form the main discussion. Because the literature detailing approaches to displaying high dynamic range images stems from linear sensors [102, 103], methods need to be adapted to deal with the output from a logarithmic camera. In section 7.2, the limitation of the dynamic range of standard displays is illustrated along with the difficulty of simple image processing algorithms in reproducing perceptibility and detail. Some current approaches to displaying high dynamic range images are also discussed, mainly exposing their inadequacies at accurately reproducing colour, image detail and scene features.

In section 7.3, tone mapping techniques that compress high dynamic range brightness information to the dynamic range of standard displays, are briefly reviewed. The special feature of brightness rather than luminance compression in most tone mapping algorithms is particularly highlighted as well its direct relevance to logarithmic sensors. This is followed by a new HDR colour representation that is suggested in section 7.4. It is proposed that by separating the luminance information from chroma, using the difference between the logarithmic responses, the chromaticity and colour ratios are maintained. This way, high dynamic range images become easier to tone map, faithfully render hue information and are most importantly, well suited for pixel responses from logarithmic imagers. It is further proposed that by using the green channel as the luminance-carrying channel, there is a net computational benefit in terms of reduced usage of processing resources for acceptable perceptual and colour quality.

In section 7.5, the differences in CIELAB error between the original image and that produced when the green channel as opposed to the usual luminance channel is tone mapped, is investigated. Using two high dynamic range images of a desktop lamp and a memorial church scene, the tone mapping operator reported by Larson *et al.* [104], was employed to both scenarios. In this investigation, the logarithm of the linearly captured image was employed to simulate the output of a logarithmic sensor. After applying the new high dynamic range colour representation, the tone mapped results are then visually and quantitatively compared to the originals using both ΔE_{76} CIELAB and ΔE^*_{76} colour error metrics. Finally, the effect of colour quantisation on tone mapping the green channel of an image are investigated; revealing a trade off between the number of bits used to represent each colour channel and the desired colour image quality for low cost applications.

7.2 Image Output

Most images from typical linear image sensors have low dynamic range and are displayed on standard TV, computer, camera or advertising screens. Often they are printed on photographic paper or other

material as a representation of the imaged scene. With the advent of HDR computer generated graphics, coupled with better architecture and improved imaging techniques in both linear and non-linear sensors, a mapping problem has arisen. This is because most contemporary displays have limited contrast ratios ranging from 100 : 1 to about 800 : 1 corresponding to just below 3 decades of brightness compared to the greater than 6 decades that some images might encompass. This is a result of the electron guns in a cathode ray tube and other LCD and TFT technologies being physically limited in creating a greater brightness range. Therefore, when high dynamic ranges are imaged, the resulting scenes are not faithfully reproduced on these media, with significant distortion and loss of detail in both the very bright and / or dark parts of the scene. A clear example is an image where both a shadow and the bright sun are portrayed; it usually lacks clarity in both the dark and bright parts of the image when displayed on typical screens; thus yielding an image bearing no resemblance to the original scene.

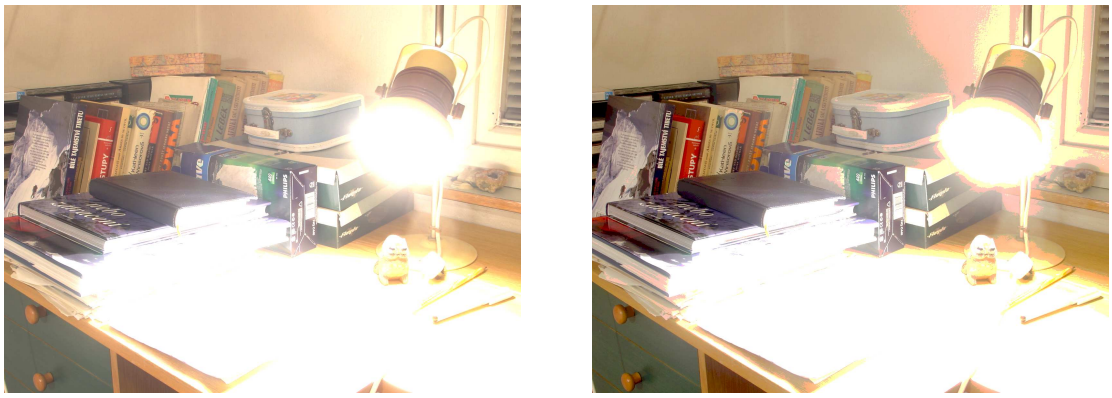

(A) (B)

Figure 7.1: A high dynamic range image, approximately 5.6 decades, displayed without any compression or extra processing - (A) - courtesy of Martin Cadik [7]. In (B), simple histogram equalisation is applied to the lamp scene creating a halo effect and undesirable perceptual artefacts.

Similarly, dyes, pigments and paints in other output media such as lithographic, painting and printing processes, can only create a white that is $20-30$ times (1.3 decades) brighter that the darkest black [102, 105]. The reason a smaller range is encountered is primarily because of the subtractive properties used to create the colours in these processes, unlike the additive ones in displays. Inevitably, printing a high dynamic range image displayed on a standard screen results in the same, if not further reduction of meaningful contrast. Figure 7.1(A), shows a high dynamic range scene illuminated by a very bright lamp. Spanning $5-6$ decades, it is clear that the bright parts of the image are over saturated while the dark areas are under exposed resulting in overall poor perceptibility of the scene[1].

[1]Note that this and the rest of the high dynamic range images in this thesis will show less perceived contrast compared to displayed version of the same image, for the reason that print media has less range than display media.

To improve the visual perception of wide dynamic range images on low contrast screens, two possible solutions exist. One approach is to radically increase the hardware capability of the low range standard displays to at least 6 decades. Ledda *et al.* [106], for instance, proposed a high dynamic range stereographic viewer that combined layered transparencies, stereo optics and strongly illuminated back light to create contrasts of up to $10,000:1$. However, in addition to suffering low resolution, this approach required that a user physically hold the HDR viewer to use and the technique had only been successfully tested using artificial static scenes. Nevertheless, better success has emerged from combining very bright low resolution LED's with high resolution low brightness LCD's to acquire contrast ratios of more than $100,000:1$ in a 16 bit display system [107–110]. Although this technology has improved to claim a $200,000:1$ brightness ratio [111, 112], this display is still far from common and is likely to remain for professional high-end use unless costs drop to comparable levels to those of current displays.

The other solution to limited range is to reduce the range of the captured image whilst preserving its perceptual quality. Such compression techniques are known to be cheaper, easier to deploy and are suitable for existing technology. Their main downside however, lies in maintaining perceptual and colour quality during compression. Simple linear mapping algorithms such as mapping the mean or maximum value of the image to that of the corresponding mean or maximum value of the display, and non-linear exponential and interval scaling have all been tried. Nevertheless, they failed to match the viewers perceptual experience of the real scene [103, 104] because of the more complex nature of vision. Even standard contrast adjusting algorithms such as histogram equalisation [3] are known to be incapable of preserving the natural features of the real scene. To illustrate, the scene in figure 7.1(A) has been histogram equalised to give figure 7.1(B). No significant improvement in perceptual quality is produced and yet visible artefacts such as the halo effect around the lamp have been created making the image even worse.

Similarly, other techniques to reduce the dynamic range of the original image such as applying power functions have not yielded better results since they lack the theoretical merit in terms of conserving colour appearance. For example, in order to examine whether displaying the logarithm of the image, rather than its linear representation, produces good rendition on standard displays, both a low and high dynamic range image were employed in this study. Figure 7.2 shows the results of this experiment where no other colour processing operation except gamma correction[2] was employed on both images. The figure shows that, despite the reduction in perceived contrast for both images, the actual colours are washed out, lack

[2]Gamma correction is a colour independent process where pixel response values are adjusted to compensate for the non-linear transfer function between the pixel intensity and the brightness of the electron guns in display systems.

clarity and are not faithfully reproduced when displayed and / or printed. This also implies that despite the unsatisfactory colour, displaying images directly from logarithmic cameras reveals details, such as the door and the electrical cord in the figure 7.2(B) and (D) respectively, that are imperceptible to linear imagers.

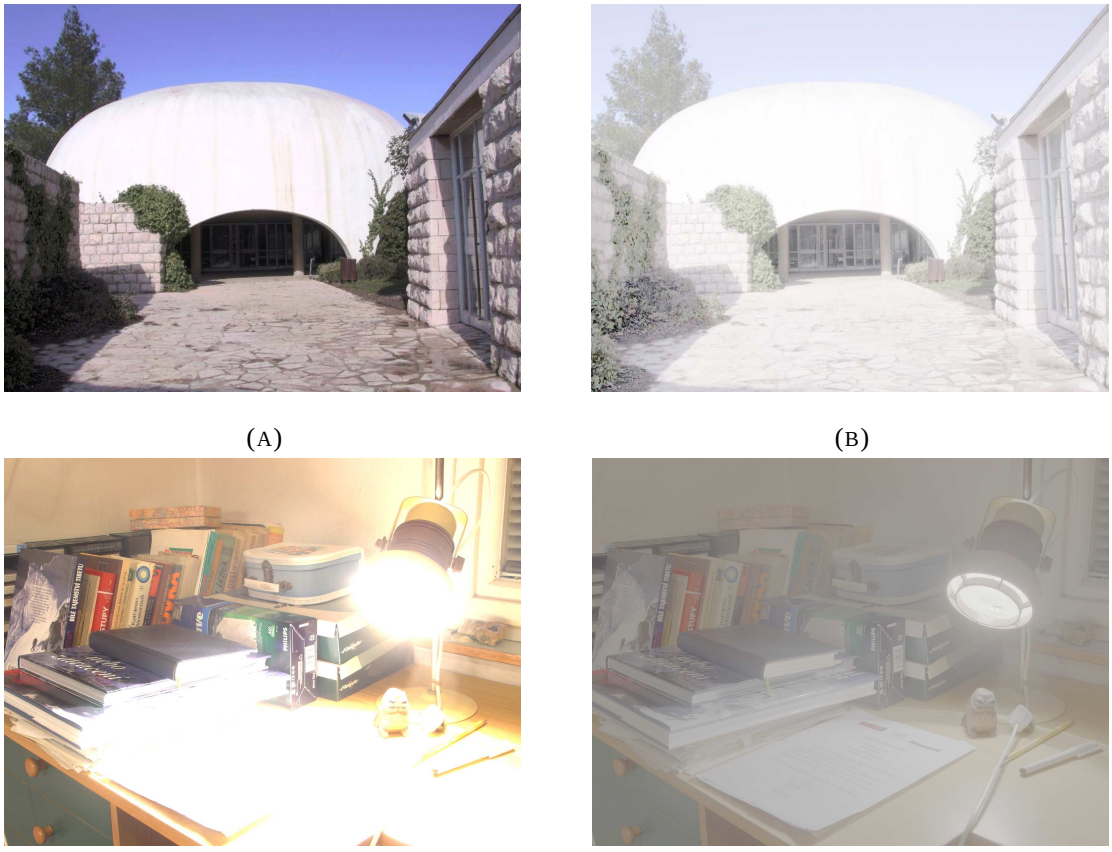

(A) (B)

(C) (D)

Figure 7.2: Examples of scenes where displaying the logarithm of an image is illustrated. (A) is a low dynamic range image of 2 decades - courtesy of Dani Laschinski [8] while (C) is a high dynamic range one of approximately 5.6 decades with the corresponding logarithm of the images shown in - (B) and (D).

Unsurprisingly, some logarithmic sensor manufacturers[3] have chosen to simply modify the image to be displayed, in a bid to reduce the dynamic range of the captured scene closer to that of a typical display. Figure 7.3 is a diagram showing logarithmic colour capture, processing and eventual display from a commercial HDRC camera. No detailed record was available for the performance of this camera in high dynamic range scenarios [1], there is an unspecified treatment of dynamic range as highlighted in the schematic. It shows each of the colour bands being modified by a different multiplicative α_k term (where

[3]Most dynamic range compression from HDR imagers are actually not addressed by manufacturers but by the users of the acquired images.

$k \in \langle R, G, B \rangle$) followed by gamma correction before display. Although it is possible these terms could be corrective coefficients for white balancing, it is likely they incorporate a factor that reduces the ranges of each of the bands to lower values pre-display - in effect, yielding results similar to the unsatisfactory images of figure 7.2.

Despite the fact that numerous image-processing operations such as spatial filtering can still be applied to different images to improve appearance and detail, acquiring a visual match with the original scene always remains a difficult task. Therefore, other approaches to rendering high dynamic range colour information on typical displays have been sought.

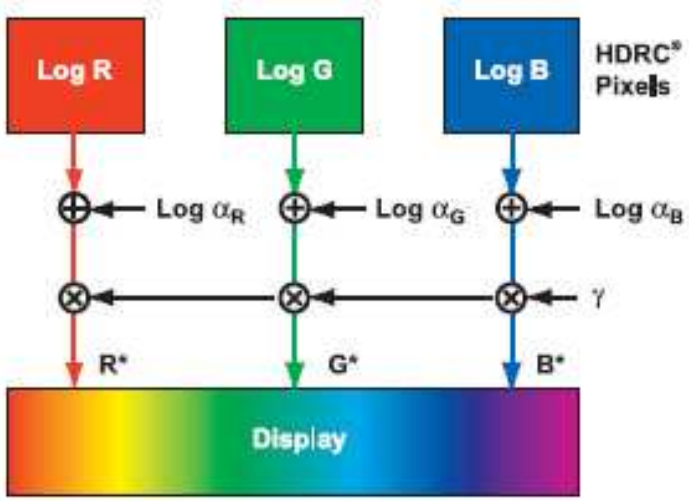

Figure 7.3: A schematic showing how colour information is manipulated and displayed by the commercial HDRC

7.3 Tone Mapping

One means by which images from high dynamic range imagers can be reproduced on standard screens is by employing compression algorithms called tone mapping or tone reproduction operators. By replicating the viewers' subjective feeling of relative brightness within different parts of a scene without actually maintaining the original brightness levels, the perception of high dynamic range in natural and artificial scenes can be reproduced on standard displays without loss of scene features, details or perceptibility. Tone mapping therefore takes advantage of the relative brightness sensitivity of the human visual system, a property discussed and employed in earlier chapters. By scaling the scene brightness levels to that of the display, the ideal tone mapping operator seeks to invoke the same visual experience of contrast

when the user looks at the scene as when it is displayed without actually using the exact values. Ideally, tone reproduction enable the completion of accurate high dynamic range colour capture, processing and display on a low dynamic range screen as shown in figure 7.4. Nevertheless, keeping all other factors constant, the level of perceptual quality of the displayed image depends on how the tone reproduction operator affects colour and other scene features during the compression of dynamic range.

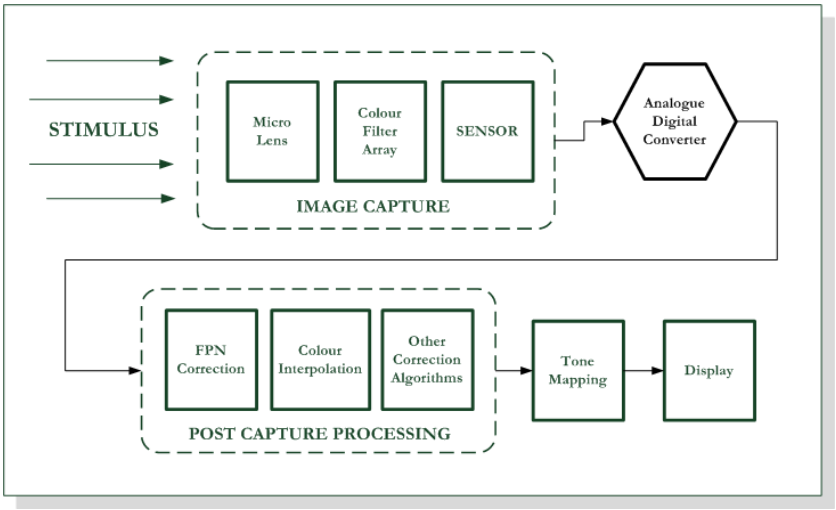

Figure 7.4: A schematic of the process of high dynamic range image capture, processing and output on standard displays.

An overview of dynamic range compression algorithms shows that tone mapping is an active area of research for several applications having many approaches that can be classified according to the way they influence luminance levels within an image. Spatially global or uniformly varying operators transform all the pixels with the same luminance in an image in the same way irrespective of their location. Conversely, spatially non-uniform or local tone reproduction operators take into account neighbourhood pixel variations and regional differences within the image [113]. These are generally more computationally complex and are better suited to images with widely varying features that warrant the extra robustness for satisfactory rendition. As might be expected, many tone mapping operators have been intricately designed to mimic human perceptual properties such as glare, colour vision, foveal adaptation and visual acuity. Although that level of detail is out of scope for this thesis, it is interesting to review the approaches some researchers have employed to underscore the depth of the tone mapping problem.

In attempting to preserve the overall impression of brightness in a displayed scene, Tumblin and Rushmeir [114] used a brightness function based on the observer's perception model of the environment to compress wide dynamic ranges to acceptable display values. Mathematical observer models of the

HVS such as light dependent visual effects, were used to compress the extreme lightning conditions in a computationally easy manner albeit only monochrome images were considered. On the other hand, Ward [115] concentrated on perceived contrast and the mapping of just noticeable difference instead of brightness. In the belief that making just discernible luminance differences in the real scene just as visible on the limited display, a scale factor could be derived using contrast sensitivity data to achieve an appropriate compression. However, because the lighting atmosphere had to be taken into account, the raw image needed to be rendered in absolute radiance units making this method not easy for common use.

Perceptual based operators have also been extended to include time-dependent factors such as light adaptation, making real-time tone mapping a possibility. For instance Ferwerda *et al.* [116] included rod and cone time-response characteristics to even out the eye's adaptation to the various light levels as the scene changes, correlating this to the just noticeable difference in his tone mapping algorithm. Pattanaik *et al.* [117] also developed an operator that computed retina-like response signals for each pixel in a scene, adding time dependent perceptual components for neural and saturation effects to improve visual adaptation at local neighbourhoods of pixel values. Algorithms like these show the complexity tone mapping operators have come to include.

One very important aspect of the tone mapping process is the way colour information is handled while trying to preserve scene perception and the fidelity of image features. Tone mapping processes primarily compress the brightness values in the high dynamic range image making the transformation essentially a gray scale operation. The separation of the luminance component from the chromatic components in the captured image then becomes essential in every tone reproduction process. In some of the earliest colour tone mapping work, Schlick [118] extracted the luminance values by obtaining the Y values using the RGB to XYZ conversion transformation. To preserve colour information, he then kept the ratios of the individual RGB colours to the luminance constant throughout the tone mapping process. Other researchers have tried modelling the rod and cone responses depending on the luminance level to capture the differences in colour reproduction [116].

Despite the multitude of various other tone reproduction techniques [119–123] and their suitability to particular applications or tasks, there is no clear industry standard tone mapping operator or method of assessment. However, the compression operator suggested by Larson *et al.* [104] has become one of the more known tone mapping operators. By exploiting the fact that the HVS adapts to a 1° foveal adaptation angle, clusters of pixels were subjected to a modified histogram equalisation hence simulating rather than maximising visibility. Interestingly, the global operator included models for glare, colour sensitivity and

glare but remained computationally simple. Consequently, it has been chosen as the operator for the investigations in this work since the results are perceptually acceptable while avoiding the more complex algorithms.

For logarithmic sensors therefore, the interest in tone mapping should only be to compress the dynamic range of the captured image while preserving visibility and colour. This is contrary to the computer graphics industry where both visibility and viewing environments are important and from where the above approaches are rooted. In addition, the process of tone mapping essentially involves manipulating the human sensation of brightness of a scene so that the same experience is obtained from a display. Considering that, brightness is determined as the logarithm of the luminance levels, logarithmic sensors readily provide these brightness values for tone mapping and therefore immediately eliminates computational burden at no extra cost. This suitability offers practical simplicity compared to linear sensors, which have the added logarithmic operation during tone mapping. Finally, the ideal solution for displaying HDR images from logarithmic sensors should be independent of the tone mapping technique if they are to be used in a vast variety of applications.

7.4 New HDR Representation

To take advantage of the suitability of logarithmic sensors for tone mapping, the imaging system can be viewed as one whole unit rather than as isolated components. Considering that good reproduction of colour depends on its accurate capture and $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ colour matching curves are most representative of any colour that can be imaged, it is possible to suggest a hypothetical colour capture scheme based on ideal XYZ filters. This would ensure the precise capture of all colours in the chromaticity diagram and avoid the inadequacies of RGB filters in current sensors[4]. Therefore, by considering a logarithmic sensor with these filters whose output is LogXYZ data, where if the differences between the LogX and LogY tristimulus coefficients and between LogZ and LogY values are taken, a new Logx'Yz' colour representation is obtained, that offers a LogY luminance channel and separate Logx' and Logz' chroma bands. Here, x' represents the ratio of the x and y chromaticities while z' is the ratio of z and y chromaticity coefficients. Using the following equations from colour trichromacy, it can be shown that this new colour representation is indeed akin to the xyY colour model. Rewriting the chromaticity equations (5.6)

$$x = \frac{X}{X+Y+Z} \qquad y = \frac{Y}{X+Y+Z} \qquad z = \frac{Z}{X+Y+Z}$$

[4]Most sensor filters are RGB filters, but because they do not accurately mimic the $\bar{r}(\lambda)$, $\bar{g}(\lambda)$, $\bar{b}(\lambda)$ functions, not all spectral colours are fully captured. Colour matrixing has to be performed post capture to correct for this deficiency.

it is possible to see that by maintaining the ratios between the tristimulus values of XYZ, which in the logarithmic domain are simple subtraction operations, the specification of the hue in the colours will remain unchanged since the ratios between the chromaticities has stayed constant as well. Because the LogY is still the luminance channel, it closely resembles the xyY colour space, which also separates the chroma and lightness, and are both derivatives of XYZ. It can then be shown that since tone mapping occurs in the luminance domain only, this colour format will not affect the specification of the hue during dynamic range compression; note that the LogY will represent brightness values readily usable for the tone reproduction operation. Hence, when the high dynamic range luminance image with luminance channel Y, is compressed to a smaller $\acute{Y}$ channel, then from equation (5.6), the x chromaticity coefficient is first given by

$$x = \frac{Y\left(\frac{X}{Y}\right)}{Y\left(\frac{X}{Y}+1+\frac{Z}{Y}\right)} = \frac{\frac{X}{Y}}{\left(\frac{X}{Y}+1+\frac{Z}{Y}\right)} \tag{7.1}$$

$$= \frac{\acute{Y}\left(\frac{X}{Y}\right)}{\acute{Y}\left(\frac{X}{Y}+1+\frac{Z}{Y}\right)} \tag{7.2}$$

and after compression,

$$\acute{x} = \frac{\acute{X}}{\acute{X}+\acute{Y}+\acute{Z}} \tag{7.3}$$

$$\text{similarly} \quad \acute{y} = \frac{\acute{Y}}{\acute{X}+\acute{Y}+\acute{Z}} \tag{7.4}$$

Hence the ratio of the energy captured by the $\overline{x}(\lambda)$ filter to the total energy, in other words the colour hue, remains the same as before tone mapping. This is the same for the other bands thus this colour representation enables an optimal way of managing tone mapping while maintaining the fidelity of colour hues. Easy tone mapping of scenes from logarithmic pixels is possible because the ratios of the chromaticities are implemented with simple subtraction operations and the luminance channel provides ready values for tone mapping.

Nevertheless, this Logx'Yz' colour format is based on idealistic $\overline{x}(\lambda)\overline{y}(\lambda)\overline{z}(\lambda)$ colour filters. RGB filters are still mainstream despite the inability of modern dyes and pigments to correctly replicate the $\overline{r}(\lambda)\overline{g}(\lambda)\overline{b}(\lambda)$ colour matching functions. In fact, current RGB filters seem to be a hybrid of ideal RGB and XYZ curves, with any subsequent colour imperfections compensated for, post capture.

For the new colour representation to be used with the predominant RGB filters, it is proposed that the new format be modified to Logr'Gb' where r' represents the ratio of the red channel to the green

colour band and b' the ratio of the blue colour band to the green [124]. This proposal is based on the fact that the human visual systems luminosity curve (figure 7.5(A)) very closely matches the green channel response - figure 5.1, and the relative brightness function is highest over the same green band - figure 7.5(B). Therefore by tone mapping the green channel, which constitutes approximately 72% of the full luminance channel, a trade-off can be made that avoids the additional calculations in determining the brightness values, LogY from LogRGB, while concurrently trying to achieve acceptable perceptual results.

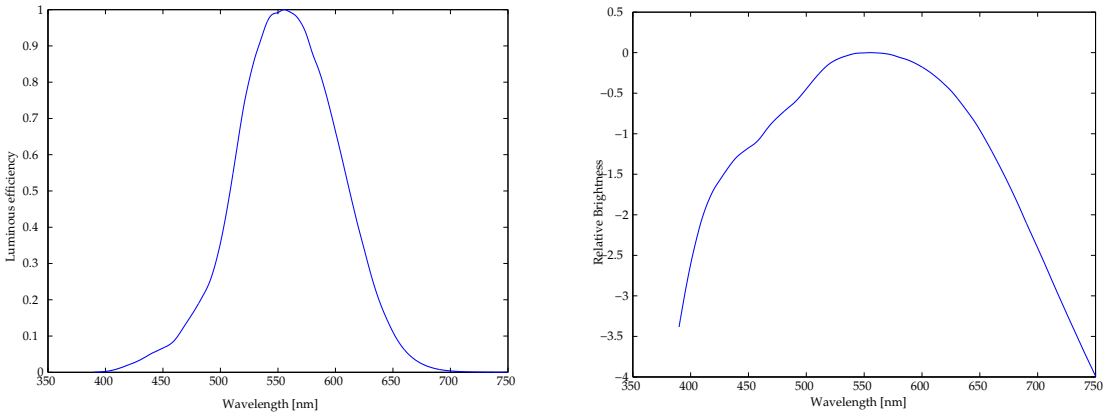

(A) (B)

Figure 7.5: Figure showing the human visual system's luminosity function in (A) and the relative brightness function (B) over the visible spectrum [9].

By keeping ratios between the red to green and blue to green constant, the colour ratios are inherently being protected before and after tone mapping hence preserving colour reproduction. This is supported by results from psychophysical experiments in the field of colour constancy that show that chromatic adaptation transforms from more stable colour ratios produce sharper and better colour images than those produced by less stable colour ratio transforms [125, 126]. Therefore, it is expected that while the chromaticity and colour hue will be the same over the tone mapping process, the lightness component of the colour definition will be altered by the compression process. In terms of visual perception, this means although there will be a shift in tone or shade of the original colour during tone mapping, collectively overall perception will be improved due to the reduced lightness range.

To show that the chromaticity is preserved when the ratios are kept constant in this dynamic range compression process, consider a colour sample $R_1G_1B_1$ such that when passed through a tone mapping process becomes $R_2G_2B_2$ where the ratios between the bands have been maintained. Then $R_2G_2B_2 = k[R_1G_1B_1]$, where k is a constant scalar. Subsequently, linear transformation to convert RGB to XYZ

can be done using a 3×3 matrix M, where Y_1 has been tone mapped to Y_2 and M represents the standard transformation matrix [98]. We can write the chromaticity coordinates x and y of the colour before tone mapping as

$$\begin{aligned} x_1 &= \frac{X_1}{X_1+Y_1+Z_1} \\ &= \frac{aR_1+bG_1+cB_1}{[R_1(a+d+g)+G_1(b+e+h)+B_1(c+f+i)]} \end{aligned}$$

and

$$\begin{aligned} y_1 &= \frac{Y_1}{X_1+Y_1+Z_1} \\ &= \frac{dR_1+eG_1+fB_1}{[R_1(a+d+g)+G_1(b+e+h)+B_1(c+f+i)]} \end{aligned}$$

where XYZ$= [M] \cdot$RGB and

$$M = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix}$$

Re-writing the chromaticity coordinates of the new colour after the mapping process gives

$$\begin{aligned} x_2 &= \frac{X_2}{X_2+Y_2+Z_2} \\ &= \frac{aR_2+bG_2+cB_2}{[R_2(a+d+g)+G_2(b+e+h)+B_2(c+f+i)]} \\ &= \frac{k[aR_1+bG_1+cB_1]}{k[R_1(a+d+g)+G_1(b+e+h)+B_1(c+f+i)]} \\ &= x_1 \end{aligned} \tag{7.5}$$

and similarly

$$\begin{aligned} y_2 &= \frac{Y_2}{X_2+Y_2+Z_2} \\ &= \frac{dR_2+eG_2+fB_2}{[R_2(a+d+g)+G_2(b+e+h)+B_2(c+f+i)]} \\ &= \frac{k[dR_1+eG_1+fB_1]}{k[R_1(a+d+g)+G_1(b+e+h)+B_1(c+f+i)]} \\ &= y_1 \end{aligned} \tag{7.6}$$

Hence, the chromaticities of both the previous and final colour after tone mapping have the same hue despite having different brightness. This Log$r'Gb'$ colour representation which preserves ratios between the colour bands, therefore offers simplicity in terms of readily available values for dynamic range mapping from a RGB filter based sensor. Figure 7.6 illustrates the computational trade-off if ideal LogXYZ filters were employed as opposed to regular RGB filters for both linear and logarithmic pixel sensors. In each scenario, the number of operational steps taken by a specific type of filter to display a captured high dynamic range image is shown along with the typical functions.

The typical steps traditional linear sensors undergo are shown on the far left of the figure, and the typical process a logarithmic sensor would undertake, in the second column. In both cases, the raw data obtained from the sensor is first corrected and routines like black level adjustment or automatic white balance carried out, before any application specific processing (for example number plate recognition). However, if a high dynamic range image is to be displayed, luminance information from the LogRGB data requires additional operations like exponentiation are needed before returning to the log domain, making the process more demanding and costlier. In the third column, use is made of the new colour representation, where the ratios between colour bands are kept constant, except that ideal XYZ filters are employed with the sole purpose of highlighting, even greater simplicity. Here there is no need for exponentiation and then the inverse, just to acquire luminance information for tone mapping. Finally in the last column, the trade off between accuracy and computation is highlighted where tone mapping the green channel is seen to be even more convenient with a step less than when using hypothetical XYZ filters. While using RGB filters makes for less computation, image quality may be compromised since the tone mapped green channel is not fully representative of the luminance inherent in the original image. However, proper tone mapping by itself changes the image appearance from its original, of which the degree of compromise, remains to be seen.

7.5 HDR Output

In order to investigate any possible compromise in quality when the green colour channel is tone mapped rather than the luminance channel, images captured using RGB filters are tone mapped using both approaches then visually compared before being studied using CIELAB error metrics. The resulting discrepancy in CIELAB coordinates when the green channel is tone mapped as opposed to the full luminance Y channel indicates the extent to which image quality may have been compromised. In this study two high dynamic range images are used; the same desk lamp image in figure 7.1(A) of approximately

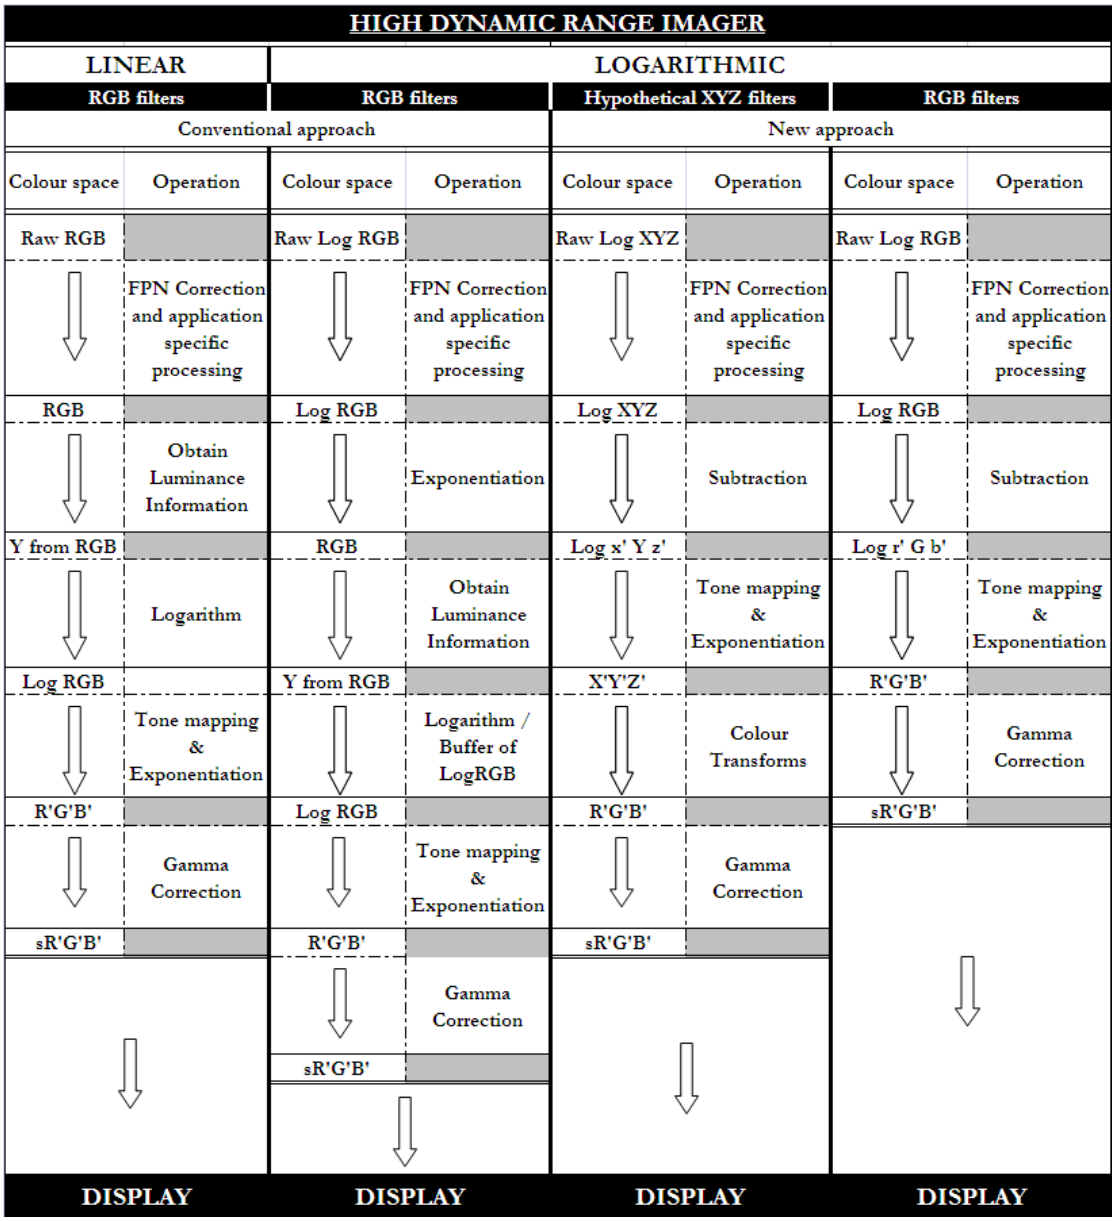

Figure 7.6: An illustration of the operational advantage of applying colour ratios for logarithmic RGB sensors on high dynamic range images and the ease of tone mapping these images.

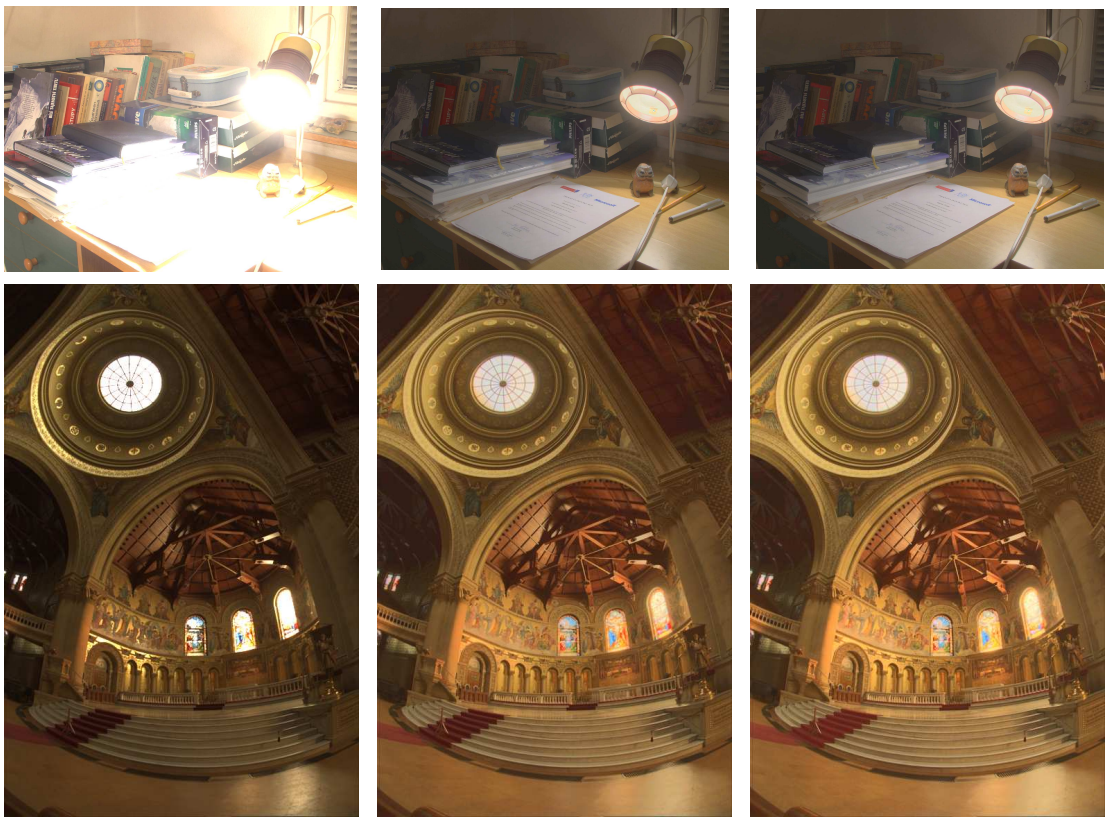

Figure 7.7: High dynamic range scenes of the desk lamp and the Stanford memorial church - courtesy of Paul Debevec [10], shown at the left, along with their tone mapped versions, the middle ones having been luminance tone mapped. The ones on the right had their green bands tone mapped.

5.6 decades dynamic range and an indoor scene of the Stanford memorial church captured using a 35*mm* SLR camera covering approximately 4.8 decade range [127]. Although these images were captured by linear cameras using multi-exposure techniques, their logarithms were used in this study to emulate what would have been captured by logarithmic sensors.

In the first scenario, the image luminance was determined from the RGB values of both images using the CIE standard conversion formula, $Y = 0.21R + 0.72G + 0.07B$, before the ratios of the individual colours bands to this luminance Y was calculated. Note that this was done following the process described in the second column of figure 7.6. Tone mapping using Larson *et al.* [104] histogram adjustment operator[5] was then performed on this Y channel after which the transformed colours were re-obtained from the constant colour ratios.

In the second scenario, the colour ratios were only determined for the red and blue channels against the green channel before tone mapping the green channel, using the same dynamic range compression

[5]In employing Larson's tone mapping operator, the models for visual acuity, glare and colour sensitivity were excluded as the prime interest in this study was replicating the compression of dynamic range.

routine as above. Again, the colours were regenerated using their ratios to give images that had a smaller dynamic range than the original scene. It can be seen from figure 7.7, that the original images on the far left show all the signs of saturation in the very bright and very dark areas expected of wide dynamic range images displayed on typical screens[6]. For example, the lamp and its surroundings are void of detail while the rafters at the top of the church are too dark to perceive meaningful features. The images in the middle column, however, have better overall visibility, acuity and clearer features. The desk contents in the upper lamp image are more vivid while the rafters in the top of the church are also clearer in the bottom image. The green tone mapped images on the right looks strikingly similar to the luminance tone mapped images. There is little to differentiate between the two as the results are amazingly similar in perceptual terms.

Figure 7.8, shows cropped portions of the two images (one being the top left corner of the lamp image and the other being the middle section of the church image) that have been tone mapped using the two methods described above. Here, both tone mapped images of the same scene were brightened by the same percentage value to create more visibility. The images on the left shows the result of tone mapping the luminance channel while the ones on the right represent an image whose green channel has been tone mapped.

Both sets of results show that the similarity between the tone mapped images is still present even at a detailed level. It is therefore difficult to perceptually differentiate the two images for any discrepancies since the colours and features appear to have been replicated equally in both cases. From an analytical viewpoint, one would expect the red hues of the original image to be most visually distorted followed by the blue colours while the green colours are expected to change least. This stems from the fact that in using the green channel for tone reproduction, the contribution of the red colour to luminance is not fairly represented as compared to say blue whose contribution to the luminance channel is less significant. Furthermore, our eyes are more sensitive to hues of red than to blue ones therefore inadequately representing both would leave the reds more conspicuous. Nevertheless, most shades appear visually matched. Although these results are far from conclusive, tone mapping the green channel appear to favour applications that do not require absolute colour reproduction or inspection but need the high dynamic range image displayed on a standard display with little effort, such as automotive (vehicular assistant) applications.

For a more scientific measure of the colour reproduction performance of a compression routine that tone maps only approximately 72% of the luminance channel, the RGB values of the compressed images

[6]In all the cases, the displayed images were gamma corrected by a factor of 2.2.

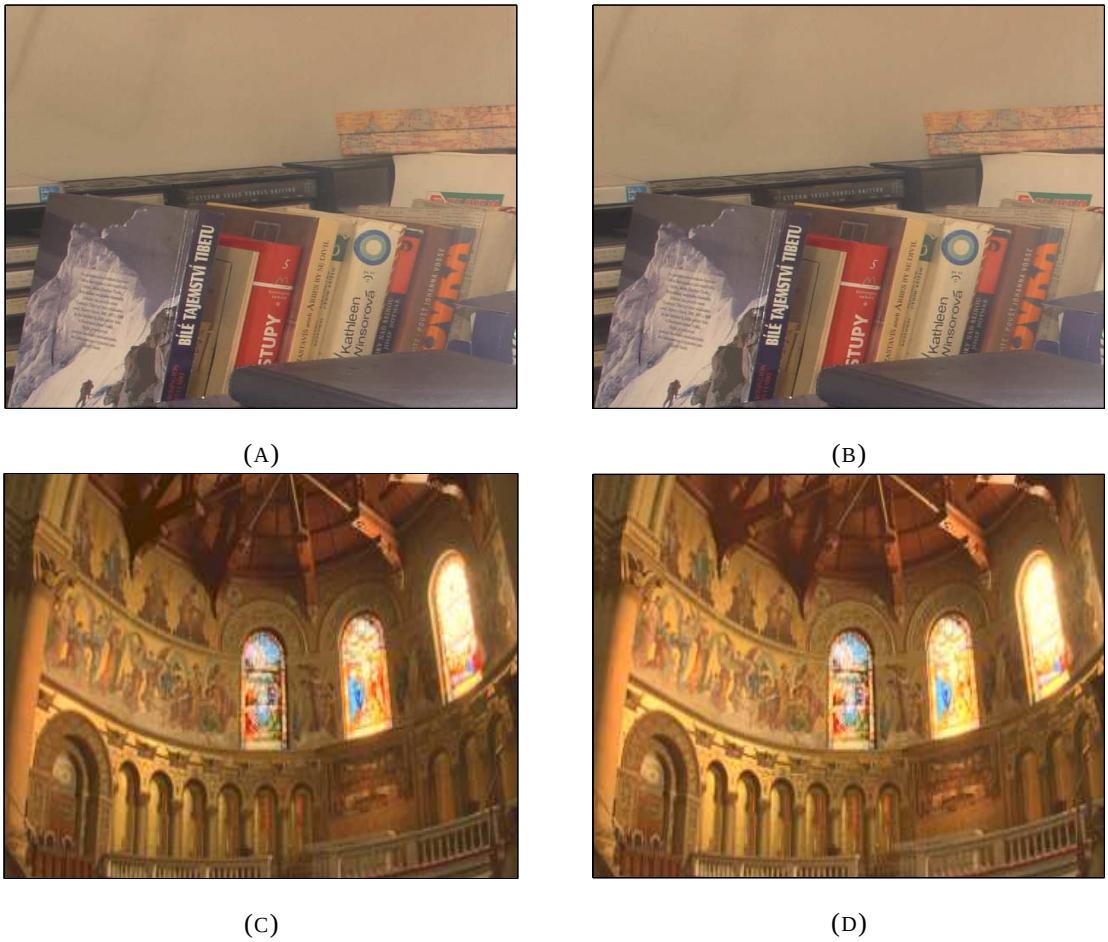

(A) (B)

(C) (D)

Figure 7.8: Figures (A) and (C) show cropped images of the lamp and the Stanford memorial church scene, luminance tone mapped while (B) and (D) show tone mapping of the green channel.

were converted to CIELAB space for perceptual colour matching against the CIELAB values of the original images. The ΔE_{76} errors of the green tone mapped images were then compared to those of the luminance tone mapped image. It should be expected that there will be high ΔE_{76} errors in this comparison since tone mapping primarily affects the scene luminance. Thus, the L component in the CIELAB error calculation is inherently changed although the location and spread of these errors should highlight any weaknesses of tone mapping the green channel as opposed to the conventional luminance channel. In fact, by showing both ΔE_{76} and ΔE^*_{76} colour and chroma error metrics respectively, the difference between the two forms of compression can be analysed even better.

In figure 7.9, the ΔE_{76} errors between the original lamp image and the luminance and green channel tone mapped versions are shown in the upper part of the image. Although the two mapped images look visually similar and the errors appear to be distributed in similar areas of the tone mapped images,

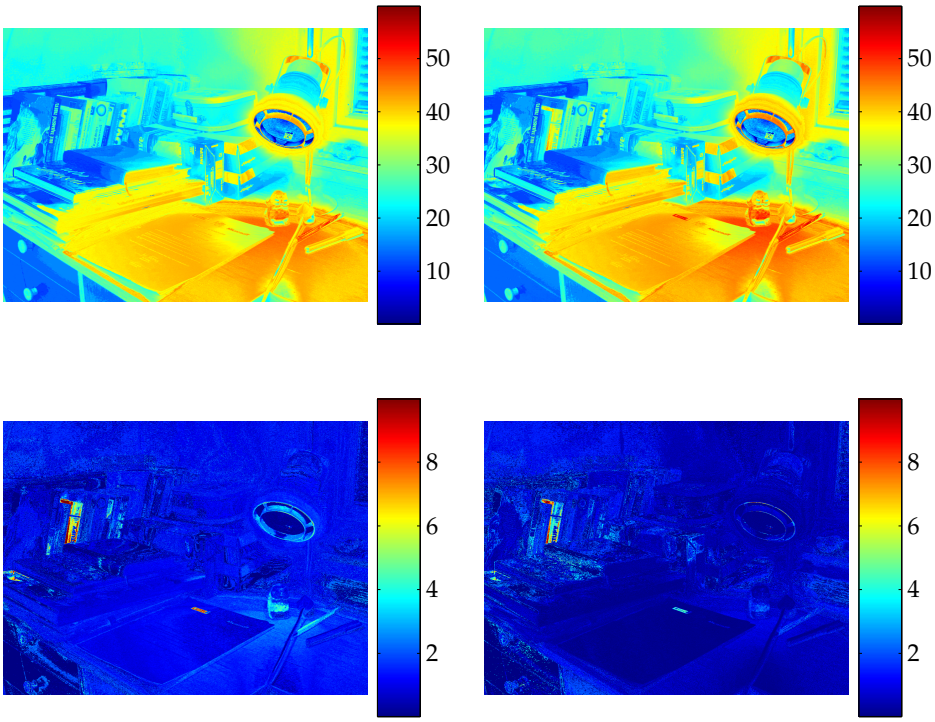

Figure 7.9: Comparing the ΔE_{76} (top images) and ΔE^*_{76} (bottom images) colour errors resulting from tone mapping the green band as opposed to the luminance channel of the lamp image. The luminance mapped images are on the left.

discrepancies appear when the ΔE^*_{76} errors are shown in the bottom part of figure. Although not fully localised to the area around the red book in the original image, this difference in error coincides with this area, for reasons discussed earlier. Still, the general error profiles remain closely matched for both tone mapped images. Similarly in the comparison of the church scene, the top half of figure 7.10, shows negligible ΔE_{76} visual difference while the bottom half highlights obvious differences in ΔE^*_{76} units when green is used to tone map the image. This time the differences spread out over the image with generally lower errors when green is used for tone mapping. The result means that quantifiable differences exist when the green channel is tone mapped but they are perceptually minute to offset any computational benefits.

A further inspection of the result of tone mapping the green channel rather than the full luminance channel can be undertaken using harsher and more obscure images. In figure 7.11, two additional uncompressed images are displayed. The left image is a church nave with windows that have the sun's rays sharply piercing through, with the rest of the image appearing dark and unclear. With a dynamic range of over 6 decades, the overall image is not well rendered, as most parts are imperceptible. Similarly, the right image shows a bright grove that appears saturated with details in the branches and the leaves

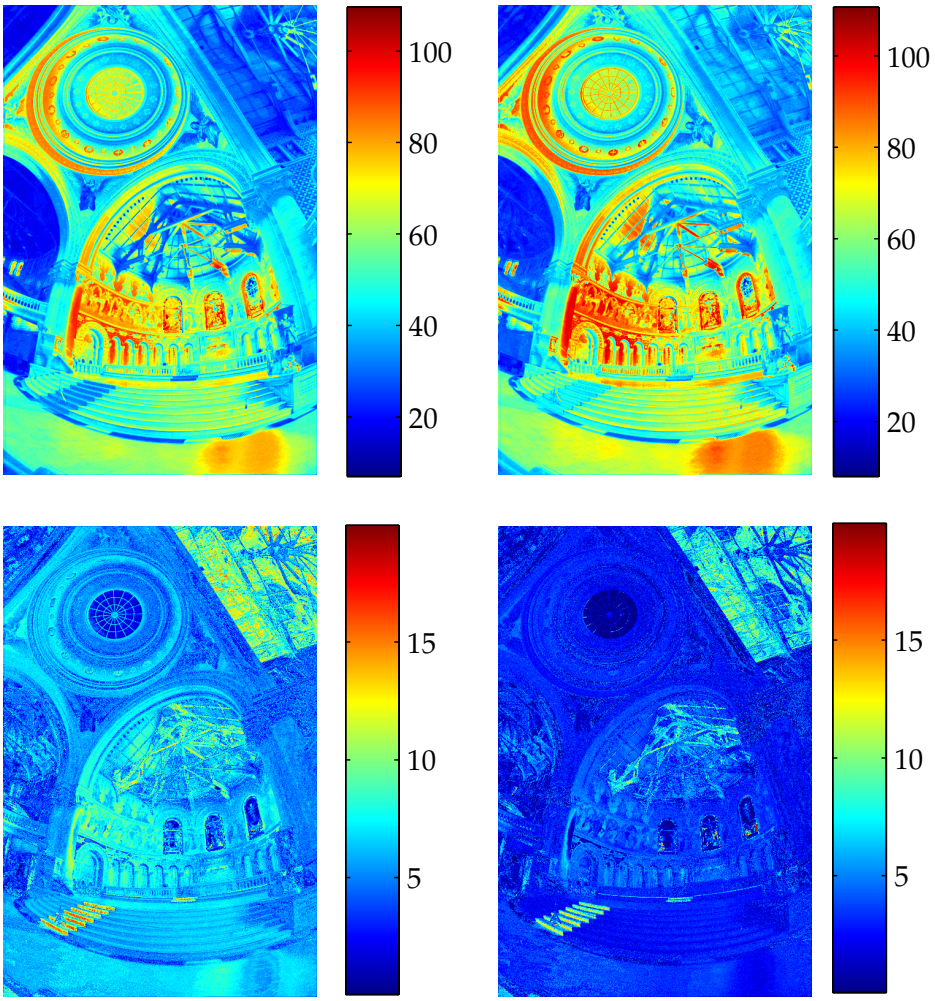

Figure 7.10: Comparing the ΔE_{76} (top images) and ΔE^*_{76} (bottom images) colour errors resulting from tone mapping the green band as opposed to the luminance channel of the church image. The luminance mapped images are on the left.

obscured.

However, after tone mapping both images, in the first instance by compressing the luminance channel and then later using only the green channel, more details become visible in parts of the image as expected. In the church nave for example, the flags and window mosaics that were previously invisible are clearer and more vivid - figure 7.12 while in figure 7.13 the leaves and smaller branches are now clearly visible. More importantly, both tone mapped versions of the two images look very similar and almost indistinguishable. As before, the ΔE^*_{76} colour errors, which should reveal any colour shifts also show little change between the luminance and green channel tone mapped images. This is more pronounced in the grove image than in the nave one possibly because there is less colour variations in the grove image

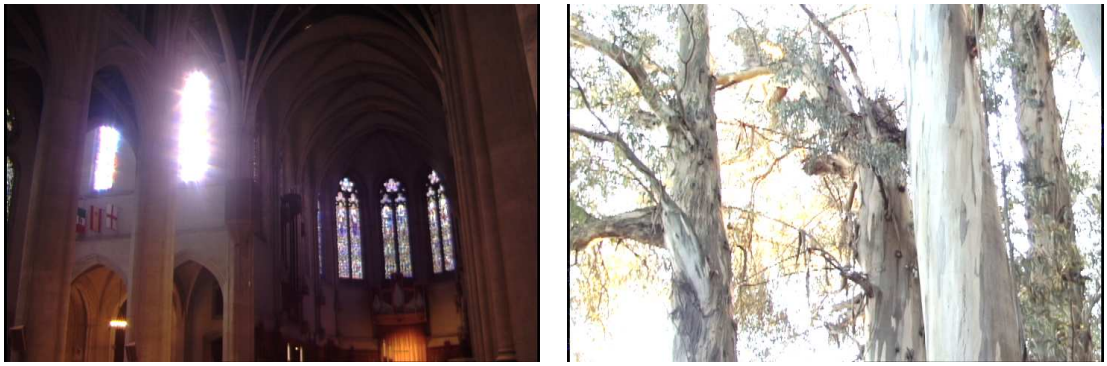

Figure 7.11: Figure showing uncompressed images of a church nave with a range exceeding 6 decades (on the left) and a grove encompassing approximately 4.4 decades. Photos are courtesy of Paul Debevec [10].

and hence less visible shifts.

Overall, the visual results match the quantitative CIELAB errors for a large part of the images leading to the conclusion that tone mapping the green channel offers an excellent trade-off of image rendition for ease of computation. It also makes for a strong argument if high dynamic range sensors are to be used in applications where high dynamic range images are to be displayed on a standard screens and not used in colour inspection applications. This specifically favours logarithmic imagers whose overall dynamic range capabilities place them at an advantage over linear sensors.

7.6 Colour Bit Length

In the previous section, it was concluded that visually similar output is more easily obtained by tone mapping the green colour channel of wide dynamic range images than with the traditional luminance data channel. Thus, the HDR output from logarithmic imagers can be suited for display on low range standard screens. These screens, typically cathode ray tubes, LCDs and TFTs, traditionally use RGB24 colour formats enabling 8 bits per colour. This has been a default standard for both the CIELAB and RGB device space, with researchers such as Stokes *et al.* [128] proving that approximately eight colour bits per channel are required for imperceptible colour differences for low dynamic range pictorial images while 10 bits per channel were required for computational precision. This is known to be representative of the extreme case where imperceptible colour differences are vital for the process in question.

However, with a moderate priority in image quality and the need to optimise computational resources in low power displays, techniques such as colour quantisation can be used to represent colour information in smaller resolution devices. This is true for colour viewfinders for modern digital cameras and other low

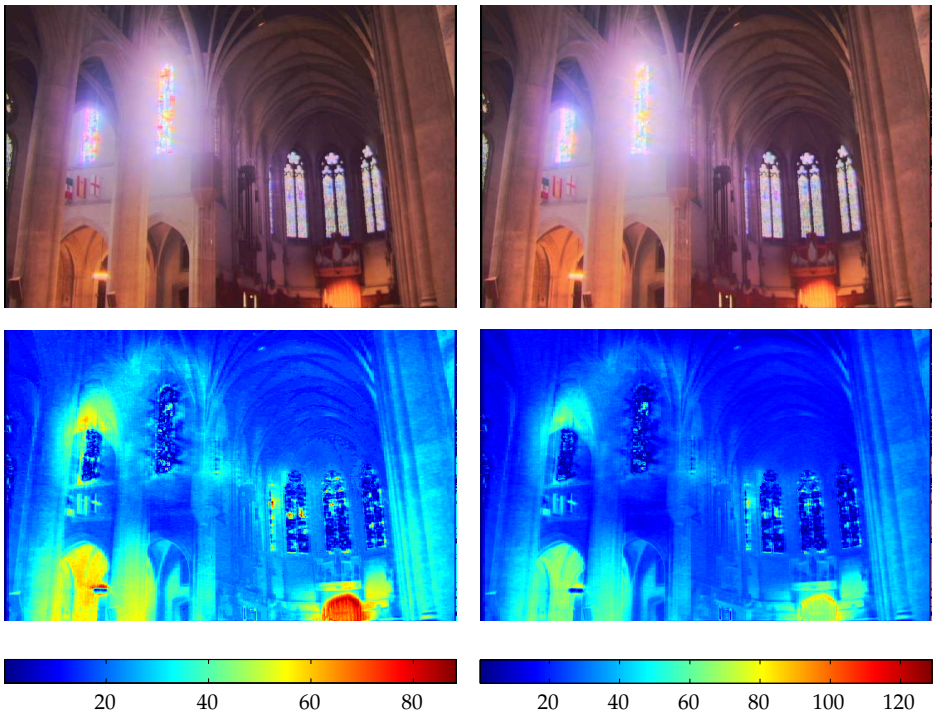

Figure 7.12: Tone mapped images of the church nave shown on the top of the figure with the luminance tone mapped one at the top left while the green tone mapped one appears at the top right. The corresponding ΔE^*_{76} errors are shown at the bottom.

cost multimedia displays, which might use logarithmic sensors as the potential imager. For instance Sharp Laboratories UK proposed a digital display with multi-format content ranging from text and graphics to video that would use only 6 bits per colour to represent images [129].

For logarithmic sensors to take advantage of such applications, it is vital that in tone mapping the green channel to compress the dynamic range of a HDR image, lower bit resolutions for low cost applications do not further degrade the quality of the displayable image. In figure 7.14, this is tested by colour quantising the results of the previous section 7.5. The original cropped image of the desk lamp is shown on the far top left before it is progressively colour quantised after tone mapping the green channel and gamma corrected. The number of bits per colour channel are gradually reduced from the original 8 to 3 and as shown, colour image quality drops gradually. As the number of bits drops below 5, the contrast in the image begins to become exaggerated since fewer shades of the colours are precisely mapped, leading to significant distortion. However, at 6 bits per colour the drop in image quality is not yet substantial for large perceptible differences from the original image. These preliminary results suggest that a simple tone-mapping algorithm may be suitable for interfacing high dynamic range cameras to low resolution LCD displays such as those found in many consumer products.

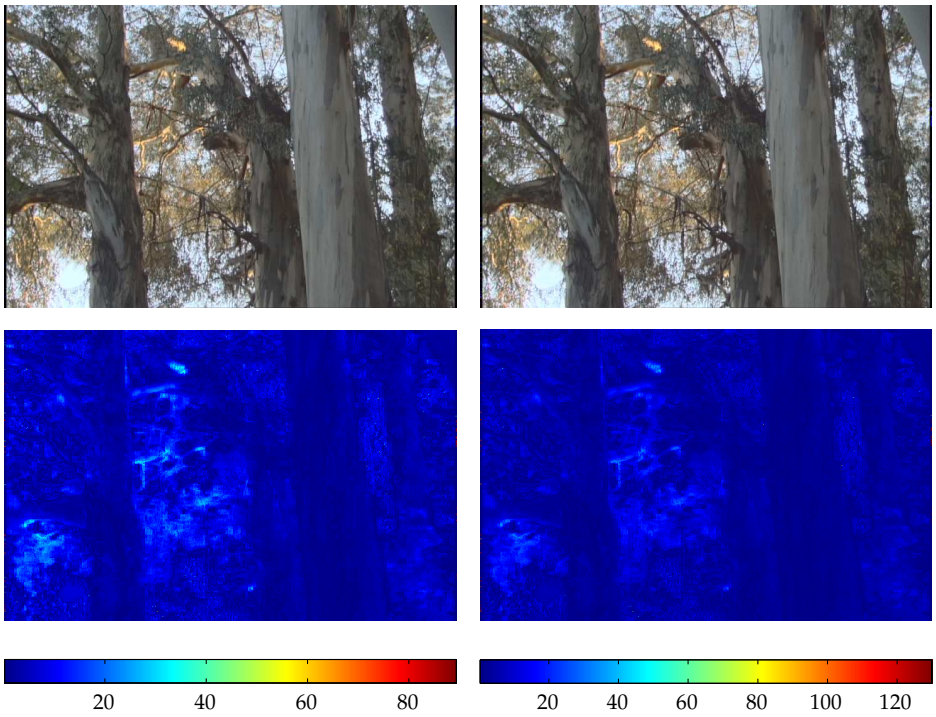

Figure 7.13: Tone mapped images of the grove shown on the top of the figure with the luminance tone mapped one at the top left while the green tone mapped one appears at the top right. The corresponding ΔE^*_{76} errors are also shown at the bottom.

7.7 Summary

In this chapter, the challenge of displaying high dynamic range images from logarithmic imagers was tackled. It was noted that although natural images and scenes from other simulated or artificial sources were capable of providing over 5 decades of illumination, standard output screens were not able to fully display this dynamic range because device physics limited output to just over 2 decades. The same problem was present for other means of output such as printing and lithographic processes, resulting from the inability of dyes and paints to create wide contrast variations.

Despite the emergence of new technologies that enable significantly higher dynamic range, they are still hugely expensive, largely untested and complex. Using techniques such as tone mapping, the computer graphics industry has managed to develop a multitude of dynamic range compression techniques that can be used to reduce the wide range present in some images to that of typical output displays. This increased visibility over both the very bright and dark areas, while enhancing scene details and features, which were previously unclear.

Since the majority of the tone mapping operators were based on linear sensors, and because the

Figure 7.14: Colour quantised, cropped image of the lamp scene whose green channel has been tone mapped. The original cropped image is shown at the top left before progressively quantising it, row wise from 7 to 3 bits per colour channel.

dynamic range compression technique essentially operates on brightness values, which are the logarithm of the luminance, the manner in which HDR images from logarithmic images were handled by tone mapping operators was re-thought. A simpler more effective but convenient colour representation was then developed, which specifically benefits logarithmic sensors during convenient and easy tone mapping while maintaining colour scene reproduction.

By taking the difference between the responses of the LogX and LogY and the LogZ and LogY components of hypothetical $\overline{x}(\lambda)$, $\overline{y}(\lambda)$, $\overline{z}(\lambda)$ filters, it was possible to simply maintain the dominant hue after tone mapping while reducing the operational steps and computational load. This was equivalent to preserving the colour ratios or the chromaticity coordinates as constants; a fact that was analytically proven. The new representation was then revised to match the fact that most filters were actually $\overline{r}(\lambda)$, $\overline{g}(\lambda)$, $\overline{b}(\lambda)$. With a Logr'Gb' colour format, where the LogG is the tone mapped channel, an even computationally easier process was expected. Thus, contrary to the popular tradition of tone mapping the full luminance, it was expected that compressing the range of the green colour would certainly introduce some colour error but possibly produce similar images in perceptual terms.

Using four high dynamic range images, both the luminance and green channels were tone mapped and the results visually checked to see if noticeable variations existed in the two approaches. It was noted

that the results in both images were very similar and almost indiscernible. A deeper study of two cropped images from both scenes showed minute visual differences even after brightening the scenes. Finally, the ΔE_{76} and ΔE^*_{76} error values between the original and the tone mapped images were compared. Inspite of the fact that the individual error values were expected to be high because of the tone mapping process, subtle differences showed in the distribution and values of the errors between the tone mapped images. It was possible to conclude that the two methods actually resulted in very close or visually similar results but which were marginally different.

Colour quantisation was then tried on a high dynamic range image after tone mapping its green channel to study the possibility of having displays for portable and low cost applications or devices, that would utilise less than 8 bit analogue to digital converters. Good perceptible results were obtained for up to about 6 bits per colour band thereby making this tone mapping approach feasible to some extent in low cost devices. It was therefore safely concluded that high dynamic range responses from a logarithmic sensors could be green-channel tone mapped, to produce strikingly good results, which are displayable in standard screens for less computational effort.

Chapter 8

Conclusion

8.1 Summary

In the current electronic imaging market, both linear and logarithmic image sensors exist. Linear sensors, composed of charged couple devices and active pixel sensors, have high signal to noise ratios offering good quality images. Both pixel designs have several merits and drawbacks making CCDs the choice for high end imaging while APS sensors have spanned many applications due to ease of integration and scalability. Although the cost of CCDs is still higher than that of CMOS APS sensors, at the expense of better image quality, they both suffer limited dynamic range, typically about 70dB, when imaging natural and commercial scenes often with greater than 120dB range.

Alternatively, logarithmic imagers, easily cover 140dB of intra-scene illumination while encoding contrast information. This way, human vision is mimicked without the need for extra circuitry and increase in costs. In addition to continuous operation needed for fast applications, the ease of tone mapping and suitability for applications like colour constancy makes them better suited for low cost high dynamic range imaging.

However, logarithmic pixel operation is stifled by low signal to noise ratios caused by significant fixed pattern noise. FPN poses the greatest challenge to the widespread use of these pixels with uncorrected FPN creating illumination misrepresentation by factors as high as 100. This thesis has discussed efforts to alleviate FPN from high dynamic logarithmic imagers in a simple but efficient manner. By using the human visual system as a reference, this work has proposed and met a very high standard of image quality after FPN correction in both monochrome and colour images.

8.1.1 Fixed Pattern Noise

The understanding of fixed pattern noise commenced by presenting an existing model for the response of a logarithmic pixel, which showed the constituents of fixed pattern noise variations in the pixels of a sensor. Apart from a spatial variation, these FPN components were shown to have significant temperature dependence hence requiring a simple and repeatable calibration process. Unlike some earlier calibration schemes that were manual or had other drawbacks of large size (reduced fill factor) or high power consumption, a self-calibrating procedure was deemed necessary especially in environments where recalibration would be frequent.

The structure of a pixel circuit the can electronically calibrated was also presented and its operation discussed. Using data from the simulated pixel circuit designed in the CADENCE circuit simulator, properties of a typical pixel from a set of $1,000$ simulated pixels were characterised as well as the impact of uncorrected fixed pattern noise.

Since previous fixed pattern noise correction routines reported their residual FPN errors in different and varying measures that have not been compared, a new perceptually based error measure was introduced in section 2.4. It was a biologically inspired metric based on the performance of the human visual system's ability to discern an illuminance. It is the first time that the FPN correction is being referenced to such as a high standard.

With the set of simulated pixels, previous offset correction methods were then benchmarked to this target contrast error measure. The disappointing results were explained using the model equation introduced in section 2.2 and residual gain variations. Although, image-processing techniques like median filtering had been suggested to reduce the remaining noise, this was deemed inadequate.

Simple offset and gain correction was then proposed to alleviate gain mismatch on the assumption that leakage currents were negligible since they only significantly affected image quality in low illumination. Using only two data points, the extraction of the offset and gain parameters was shown and the correction for these forms of mismatch demonstrated at any photocurrent. However, because of changing sensitivity along the photocurrent range and the need to attain the target contrast, the choice of these two data points had to be carefully selected. At the same time, numerical errors had to be minimised while covering a wide dynamic range. Consequently, the two photocurrents had to be as different as possible with the smaller current being at least 100 times larger than the c_j parameter. In empirically determining the larger photocurrent, it was found that there was a higher error as the photocurrent moved further from the smaller photocurrent. Nevertheless, errors of about 5% were realised. This was lower than the ones

reported in earlier studies in table 2.1.

Simple offset and gain correction was also benchmarked along with the best offset and gain correction method - the mean sum square optimisation method. It was found that the two methods were similar over the same range and despite the absolute contrast error not improving at high currents, the simple offset and gain method was simpler, more convenient, and less computational making it a more viable choice for fixed pattern noise correction.

8.1.2 Model Derivation & Parameter Extraction

It was noted that just as in the low illumination region of pixel operation, contrast errors were unexpectedly as high as at the high end of the current spectrum. This was investigated and, as suspected, the logarithmic transfer function becomes invalid in this region since the load transistor is being driven into moderate inversion. A new square law relationship was shown to exist between the drain source current I_{DS} of the load transistor and its gate source voltage. Therefore, modelling the response of a pixel using a logarithmic transfer function is invalid.

To model logarithmic pixel response over the entire photocurrent region, a new four parameter model based on the EKV MOS transistor model was developed that smoothly represents pixel behaviour over eleven decades of photo stimulus. The model was found to make an accurate fit with the simulated pixel's response when the four parameters are extracted using a regressive non-linear algorithm in MATLAB. Although the model errors from the fit were all below 3.5mV, this method of extracting the parameter was computationally demanding and thought to provide a burden for real-time applications.

In order to reduce the effort and time consumed in the parameter extraction procedure, a different approach was proposed. The method which focussed on specific regions of the photocurrent range, helped reduce the four parameter model to a simpler form that could be used in the extraction of the coefficients. Five conveniently chosen data points were used in this procedure, resulting in a similar fit to the function minimisation method, only for higher error at the high end, peaking at about 6mV. Although this error was at least two fold in comparison to the optimisation result, there was a tremendous improvement in computation time and effort. Principally, the simple extraction technique provided ease, convenience and accuracy over a wide range that was not possible previously. However, because of the impracticality of measuring currents in real circuits required in the simple extraction approach, coupled with the need to retain the logarithmic properties of pixels, the optimised two parameter logarithmic model was considered prime in the development of a practical fixed pattern noise correction technique.

In this quest, two approaches were discussed. In the first procedure, all pixel responses were corrected

to the average pixel response at the point of calibration, the assumption being that all corrected pixels constituted a mean offset and mean gain. This was done using pixel responses only. In the second method, the responses were also corrected for offset and gain variations without the actual knowledge of photocurrents but by taking the ratio of two offset corrected terms. In both methods, the significance of the location of the two calibration points was retained, resulting in easy and convenient strategies for FPN infested image correction. Lastly, the operational steps in a practical calibration routine were outlined thereby emphasising the simplicity and ease of these approaches to FPN correction.

8.1.3 Monochrome FPN Correction

The contrast error measure, previously used for benchmarking and characterising fixed pattern noise, was adjusted to mimic the way the human visual system adapts to perceive non-uniformity. In particular, the eye's higher sensitivity to relative illumination changes rather than to absolute ones, was central to the adoption of a more relevant relative contrast error measure. Using simulation data, the absolute contrast was shown to have a uniformly low standard deviation over the entire current range, which implied that contrast error was systematic. Using this relative contrast error, simulation data easily achieved the contrast target for 7 decades and was only limited by the level of the circuit simulator. It was also proved that relative contrast error could be easily determined from practical electronic image correction techniques mentioned in section 3.5 making it a pertinent and viable error measure.

Although, FPN correction on simulation data, resulted in achieving the contrast target over a wide range, it remained unproven with real pixel data. For this reason, the practical image correction techniques were tested on pixel responses from a 200×100 array. Before this was performed, the array was characterised to investigate the distribution and the sources of the various forms of non-uniformity. With typical behaviour from the pixels, column to column FPN was realised to be greater than row-to-row FPN, due to the extra transistors in the column calibration circuitry. This was confirmed by maps of extracted offset and gain parameters, which also highlighted the dominance of offset mismatch as well as the significance of correcting for gain mismatch.

The actual correction of fixed pattern noise was shown using the two methods; namely the mean response method and the photocurrent ratios method. A relative contrast error target of 1% was achieved over 4 decades of illumination and an extra decade at 2% contrast. This higher error was attributed to higher c_j and d_j parameter variations in the pixel sensor and to temporal noise which, despite being reduced by averaging, was still considerable.

The effects of the analogue to digital converter were also emulated by adding quantisation error to

the pixel responses and correcting for FPN. This was to make a more accurate assessment of the image quality from a HDR imager. The choice of quantisation level was explained and applied to the real pixel responses resulting in an increase in relative contrast error and the consequential decrease of dynamic range under 1% accuracy. The level of quantisation was also shown to affect the number of bits with higher accuracy requiring a longer bit length. This was found to suit the variety of applications that have different accuracy and dynamic range needs.

Conclusively, a high dynamic range monochromatic logarithmic sensor can be easily calibrated using an electronic procedure. Wide dynamic range monochrome or gray scale images can be corrected for offset and gain variations in a fast and accurate way; this accuracy being about as good as the human visual system can detect over 4 decades or at least 2% contrast sensitivity over 5 decade scenes.

8.1.4 Colour

The foundations for the assessment of high quality fixed pattern noise correction of colour images from a logarithmic sensor were laid beginning with basic colour theory. In addition, the means of capturing or sensing colour information using colour filter arrays were also detailed as part of the sensor colour processing chain. Along with the numerous ways in which standard colours are represented, the equations that defined the most commonly used RGB & XYZ colour models were also highlighted. After discussing the merits and drawbacks of these colour spaces, the lack of perceptual linearity and intuitiveness in colour representation were confirmed using Macadam's ellipses on the chromaticity diagram.

The CIELAB colour space was then discussed as a standard perceptually uniform, device independent space from which many colour difference formulae were based and on which most of the colour work in this thesis was focused. Despite the suggested alternatives to the ΔE_{76} error metric, its use in various industries and applications, made it the most widely used and quoted metric in the colour imaging field. Consequently, it was adopted for this work. As a measure of acceptable visual quality, less than 3 ΔE_{76} units was generally recommended for logarithmic sensors, a value in agreement with research in the art and visual science world.

Munsell colours were introduced as a valid and authentic tool set, from which logical deductions could be drawn after analytical simulations had been performed. This is because they sufficiently covered the colour gamuts of typical displays and had been used in various research studies. Insights were inferred about the dynamic range of colour, ratios of the colour components and how perceptually incorrect certain values of be ΔE_{76} error were from others. Overall, definitions of a pertinent colour error metric and a data set, fit for analysing the quality of FPN correction in high dynamic range colour images were developed.

8.1.5 Colour FPN Correction

The results of correcting fixed pattern noise in monochrome images were extended to high dynamic range colour scenes. Because these images compose the vast majority of user images, this study was deemed significant if logarithmic sensors are to be useful in many high dynamic range-imaging applications.

Using basic but essential knowledge of colour reproduction from chapter 5, the link between contrast error and colour quality was first investigated. By studying the effect of absolute error in each of the constituent colour bands, the limitations of model error on the overall perception of colour was deemed tolerable. This led to a MATLAB simulation using the standard set of Munsell colours over a 10 decade range of photo-illumination.

A white reference colour was first added to the Munsell colour set as a way of transforming RGB to the perceptual CIELAB, but also as a means of normalising all the colours at any illumination. Using typical values of the a_j, b_j, c_j and d_j parameters, the response of a pixel over the illumination range, allowed for the extraction of offset and gain from the two parameter model. Modelling for colour assumed all the colours imaged the pixel one at time at each illumination. Thus, using the extracted offset and gain parameters, the estimated red, green and blue photocurrents were obtained for each illumination. After referencing all the colours to the reference white at that illumination, the estimated RGB colours were converted to CIELAB before comparing to their original RGB colour values.

By investigating the shifts in the positions of the colours on a chromaticity diagram at different illuminations, it was possible to begin to determine the effects of absolute error on colour. Fewer chroma shifts were observed in the purely logarithmic region while worse ones were present outside this region. CIELAB values confirmed this trend with colour error falling below three ΔE_{76} or ΔE_{00} units for most of the photocurrent range. An analytical investigation into the low CIELAB values then confirmed earlier predictions in section 6.2. Thus, residual absolute errors were actually manageable since they do not adversely affect colour perception.

However, to mimic real sensor noise conditions, typical fixed pattern noise variations were added to the simulation of over $1,000$ pixels. Results revealed that the addition of FPN and subsequent image correction did not change the perception of colours over the entire range. Then, ADC quantisation noise effects were also added to the MATLAB simulation to match the typical imaging process even closer. Digitisation noise was therefore added to the previous simulations for the same $1,000$ pixels. As the step size was increased, a corresponding drop in colour accuracy was realised. Inspite of this, a 5 decade range was still obtainable at less than 3 ΔE_{76} units of colour accuracy for a 10 bit ADC.

Overall, this work revealed the very high potential for high dynamic range monochrome and colour images from a logarithmic sensor that performs simple fixed pattern noise correction. The resulting images are perceptually indistinguishable over at least 5 decades making them ideal for most contemporary high dynamic range image applications.

8.1.6 HDR Logarithmic Output

The challenge of displaying high dynamic range images from logarithmic imagers was tackled. It was noted that although natural images and scenes from other simulated or artificial sources were capable of providing over 5 decades of illumination, standard output screens were not able to fully display this dynamic range because device physics limited output to just over 2 decades. The same problem was present for other means of output such as printing and lithographic processes, resulting from the inability of dyes and paints to create wide contrast variations.

Despite the emergence of new technologies that enable significantly higher dynamic range, they were still hugely expensive, widely untested and complex. Using techniques such as tone mapping, the computer graphics industry has managed to develop a multitude of dynamic range compression techniques that can be used to reduce the wide range present in some images to that of typical output displays. This increased visibility over both the very bright and dark areas while enhancing scene details and features, which were previously unclear.

Since the majority of the tone mapping operators are based on linear sensors and because the dynamic range compression technique essentially operates on brightness values, which are the logarithm of the luminance, the way HDR images from logarithmic images were handled by tone mapping operators was re-thought. A simpler more effective but convenient colour representation was then developed, which specifically benefits logarithmic sensors during convenient and easy tone mapping while maintaining colour scene reproduction.

By taking the difference between the responses of the LogX and LogY and the LogZ and LogY components of hypothetical $\overline{x}(\lambda)$, $\overline{y}(\lambda)$, $\overline{z}(\lambda)$ filters, it was possible to simply maintain the dominant hue after tone mapping while reducing operational steps. This was equivalent to preserving the colour ratios or the chromaticity coordinates as constants; a fact that was analytically proven. The new representation was then revised to match the fact that most filters were actually $\overline{r}(\lambda)$, $\overline{g}(\lambda)$, $\overline{b}(\lambda)$. With a Logr'Gb' colour format, where the LogG is the tone mapped channel, an even computationally easier process was demonstrated.

Using four high dynamic range images, both the luminance and green channels were tone mapped

and the results visually checked to see if noticeable variations existed in the two approaches. It was noted that the results in both images were very similar and almost indiscernible. A deeper study of two cropped images from both scenes showed minute visual differences even after brightening the scenes. Finally, the ΔE_{76} and ΔE^*_{76} error values between the original and the two tone mapped images were compared. Inspite of the fact that the individual error values were expected to be high because of the tone mapping process, subtle differences showed in the distribution and values of the errors between the tone mapped images. It was possible to conclude that the two methods actually resulted in very close or visually similar results but which were marginally different.

Colour quantisation was then performed on a high dynamic range image after tone mapping its green channel to study the possibility of having displays for portable and low cost applications or devices, that would utilise less than 8 bit analogue to digital converters. Good perceptible results were obtained for up to about 6 bits per colour band thereby making this tone mapping approach feasible in low cost devices. It was therefore safely concluded that high dynamic range responses from a logarithmic sensors could be green-channel tone mapped, to produce strikingly good results, which are displayable in standard screens for less computational effort.

8.2 Future Work

8.2.1 Colour Results

The effect of the two parameter correction method on monochrome images captured by a logarithmic sensor were first analysed using simulation data in chapters 2 and 3. Using a real sensor, data from a 200×100 pixel array was then employed to validate the analytical results. However, when colour images were studied, only simulated results were used. The lack of a colour logarithmic sensor was the main difficulty in obtaining real data to confirm the simulation results of chapter 5 and 6.

As part of any future final justifications for the feasibility of applying the two parameter model as a correction routine in colour images, it is pertinent that the simulation results be validated using real data. Using a sensor designed to include the electronic calibration scheme mentioned in chapter 8.1.1 and with overlaid colour filter arrays, a wide range of photocurrents can be achieved where the response of each pixel is in fact the response of only one photoreceptor class. As with monochrome data from the 200×100 array, the parameters can be extracted to give the estimated red, green and blue photocurrents from two carefully chosen currents. Similarly, these currents can be obtained from the CADENCE circuit

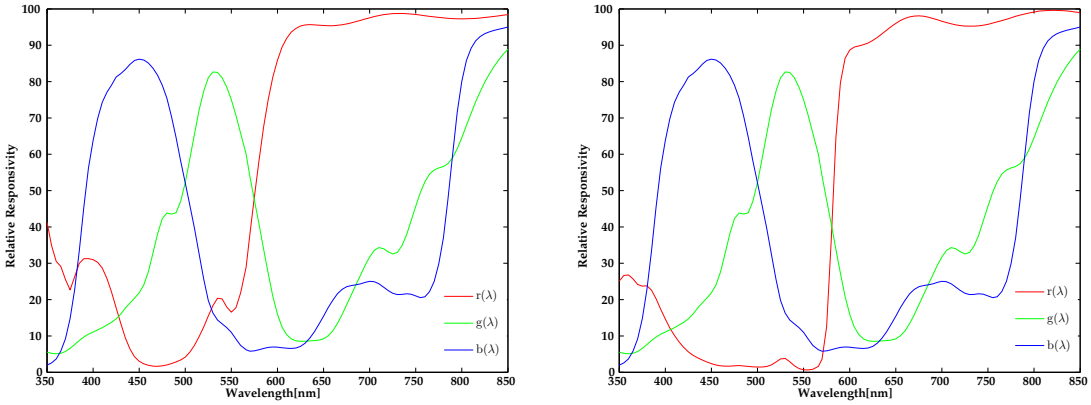

(A) (B)

Figure 8.1: (A) and (B) show examples of the relative spectral responsivities of 1.1μm thick RGB filters in the X14 and XIBIS5 sensors respectively [11].

simulator since measuring photocurrents from a sensor is difficult.

A demosaicing algorithm can then be used to obtain the complete red, green and blue photocurrents for each pixel using the relevant values from its neighbours. The CIELAB values for each pixel can be calculated before and after fixed pattern noise correction allowing for a comparison of colour error across the entire range.

Lastly, a visual confirmation of the results of this correction can be achieved if an entire camera system was built from which real monochrome and colour images can be taken. The ability to focus a high dynamic range colour image and then correct for any mismatch using either of the means mentioned at the end of chapter 8.1.2 would offer the chance to visually validate the correction routine.

8.2.2 Colour Matrixing

One feature that real colour sensors, such as one that might be required to validate previous colour work possess, is the failure to accurately capture colour information. The implicit assumption in all the simulations in this work was that the RGB filters on the colour filter arrays were ideal with smooth cut-offs at different wavelengths for the respective bands. This is not valid for either linear or logarithmic sensors. Figure 8.1, shows typical spectral characteristics of two industrial logarithmic imagers. The curves are clearly far from the $\bar{r}(\lambda)$, $\bar{g}(\lambda)$, $\bar{b}(\lambda)$ or the $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ that were discussed in section 5.2.1. Neither the wavelength nor relative sensitivities of the colour matching functions are accurately reproduced by these sensors.

Due to the non-ideal nature of this and other sensors in the market, any colours that are captured by

these imagers will always produce RGB component values that are different from the ideal RGB values of the scene being captured. This is true despite correcting for all types of noise in the capturing process.

One approach that will have to be employed to correct for this source of colour error in any future work, is colour matrixing - applied immediately after colour interpolation/demosaicing - figure 5.2. In this procedure, colour filter overlap is compensated for by comparing the RGB values of every patch of a standard Macbeth colour checker, to those obtained after imaging the same patch using a real sensor under the same illumination [41, 75]. By optimising the elements of a 3×3 correction matrix, perceptual colour from the sensor is brought to parity with that of standard colour charts. This approach is no different from what current linear and logarithmic sensors use to correct for inherent colour shifts during capture.

8.2.3 Temporal Noise

Logarithmic sensors are known to have voltage swings in the order of a few hundreds. For instance for a typical 50mV/decade sensor, imaging a six decade range would mean about 300mV is realised. With additional FPN variations, this is smaller than the voltage swings of about 1.5V enjoyed by some linear sensors. This makes logarithmic sensors particularly sensitive, to other noise sources apart from quantisation and fixed pattern noise.

The procedure to correct for fixed pattern noise that was discussed in this work included quantisation noise in its analysis. It also yielded excellent results that make noise sources like temporal noise increasingly significant. Composed mainly of readout noise and thermal noise, it was minimised by averaging pixel responses from the 200×100 sensor array. Each pixel was read eight times reducing the amount of temporal noise by a factor of $\sqrt{8}$. This means for a 1mV peak-to-peak temporal noise measurement, the residual noise level after averaging is approximately 0.1768mV rms - a value equivalent to just less than 1% contrast accuracy for a 50mV/decade pixel. Although averaging has tremendously reduced its significance, this is not possible in real sensors.

Since reading a pixel numerous times slows down performance and might need to be repeated so many times to achieve a significant effect, the averaging of pixel responses is particularly cumbersome and time consuming for large arrays. It becomes increasingly impossible if the scene also changes continuously - such as in automotive and security applications - consequently the reduction of the effects of temporal noise is being addressed from a design rather than processing perspective. Ideas like increasing the size of the calibration sources to reduce variability and incorporating substrate level ADC's have been floated [78] with some level of success. Nevertheless, the effect of temporal noise on logarithmic sensor

performance could be studied even further for a better understanding of its impact on logarithmic HDR imaging.

Appendix A

Munsell Colour Nomenclature

Munsell Colour Nomenclature					
Index	RED	Index	RED	Index	RED
1	2.5R 9/2	46	5R 2.5/2	91	7.5R 7/8
2	2.5R 8/2	47	5R 8/4	92	7.5R 6/8
3	2.5R 7/2	48	5R 7/4	93	7.5R 5/8
4	2.5R 6/2	49	5R 6/4	94	7.5R 4/8
5	2.5R 5/2	50	5R 5/4	95	7.5R 7/10
6	2.5R 4/2	51	5R 4/4	96	7.5R 6/10
7	2.5R 3/2	52	5R 3/4	97	7.5R 5/10
8	2.5R 2.5/2	53	5R 7/6	98	7.5R 4/10
9	2.5R 8/4	54	5R 6/6	99	7.5R 6/12
10	2.5R 7/4	55	5R 5/6	100	7.5R 5/12
11	2.5R 6/4	56	5R 4/6	101	7.5R 4/12
12	2.5R 5/4	57	5R 3/6	102	10R 9/1
13	2.5R 4/4	58	5R 7/8	103	10R 8/1
14	2.5R 3/4	59	5R 6/8	104	10R 7/1
15	2.5R 7/6	60	5R 5/8	105	10R 6/1
16	2.5R 6/6	61	5R 4/8	106	10R 5/1
17	2.5R 5/6	62	5R 7/10	107	10R 4/1
18	2.5R 4/6	63	5R 6/10	108	10R 3/1
19	2.5R 3/6	64	5R 5/10	109	10R 2.5/1
20	2.5R 7/8	65	5R 4/10	110	10R 9/2
21	2.5R 6/8	66	5R 6/12	111	10R 8/2
22	2.5R 5/8	67	5R 5/12	112	10R 7/2
23	2.5R 4/8	68	5R 4/12	113	10R 6/2
24	2.5R 7/10	69	5R 6/14	114	10R 5/2
25	2.5R 6/10	70	5R 5/14	115	10R 4/2
26	2.5R 5/10	71	5R 4/14	116	10R 3/2
27	2.5R 4/10	72	7.5R 9/2	117	10R 2.5/2
28	2.5R 6/12	73	7.5R 8/2	118	10R 8/4
29	2.5R 5/12	74	7.5R 7/2	119	10R 7/4
30	2.5R 4/12	75	7.5R 6/2	120	10R 6/4
31	5R 9/1	76	7.5R 5/2	121	10R 5/4
32	5R 8/1	77	7.5R 4/2	122	10R 4/4
33	5R 7/1	78	7.5R 3/2	123	10R 3/4
34	5R 6/1	79	7.5R 2.5/2	124	10R 8/6
35	5R 5/1	80	7.5R 8/4	125	10R 7/6
36	5R 4/1	81	7.5R 7/4	126	10R 6/6
37	5R 3/1	82	7.5R 6/4	127	10R 5/6
38	5R 2.5/1	83	7.5R 5/4	128	10R 4/6
39	5R 9/2	84	7.5R 4/4	129	10R 8/8
40	5R 8/2	85	7.5R 3/4	130	10R 7/8
41	5R 7/2	86	7.5R 7/6	131	10R 6/8
42	5R 6/2	87	7.5R 6/6	132	10R 5/8
43	5R 5/2	88	7.5R 5/6	133	10R 4/8
44	5R 4/2	89	7.5R 4/6	134	10R 7/10
45	5R 3/2	90	7.5R 3/6	135	10R 6/10

Munsell Colour Nomenclature					
Index	RED	Index	YELLOW RED	Index	YELLOW RED
136	10R 5/10	179	5YR 7/2	224	7.5YR 5/8
137	10R 7/12	180	5YR 6/2	225	7.5YR 8/10
138	10R 6/12	181	5YR 5/2	226	7.5YR 7/10
139	10R 5/12	182	5YR 4/2	227	7.5YR 6/10
		183	5YR 3/2	228	7.5YR 7/12
Index	YELLOW RED	184	5YR 9/4	229	10YR 9/1
140	2.5YR 9/2	185	5YR 8/4	230	10YR 8/1
141	2.5YR 8/2	186	5YR 7/4	231	10YR 7/1
142	2.5YR 7/2	187	5YR 6/4	232	10YR 6/1
143	2.5YR 6/2	188	5YR 5/4	233	10YR 5/1
144	2.5YR 5/2	189	5YR 4/4	234	10YR 4/1
145	2.5YR 4/2	190	5YR 8/6	235	10YR 3/1
146	2.5YR 3/2	191	5YR 7/6	236	10YR 2.5/1
147	2.5YR 2.5/2	192	5YR 6/6	237	10YR 9/2
148	2.5YR 8/4	193	5YR 5/6	238	10YR 8/2
149	2.5YR 7/4	194	5YR 4/6	239	10YR 7/2
150	2.5YR 6/4	195	5YR 8/8	240	10YR 6/2
151	2.5YR 5/4	196	5YR 7/8	241	10YR 5/2
152	2.5YR 4/4	197	5YR 6/8	242	10YR 4/2
153	2.5YR 8/6	198	5YR 5/8	243	10YR 3/2
154	2.5YR 7/6	199	5YR 7/10	244	10YR 9/4
155	2.5YR 6/6	200	5YR 6/10	245	10YR 8/4
156	2.5YR 5/6	201	5YR 7/12	246	10YR 7/4
157	2.5YR 4/6	202	5YR 6/12	247	10YR 6/4
158	2.5YR 8/8	203	7.5YR 9/2	248	10YR 5/4
159	2.5YR 7/8	204	7.5YR 8/2	249	10YR 4/4
160	2.5YR 6/8	205	7.5YR 7/2	250	10YR 8/6
161	2.5YR 5/8	206	7.5YR 6/2	251	10YR 7/6
162	2.5YR 4/8	207	7.5YR 5/2	252	10YR 6/6
163	2.5YR 7/10	208	7.5YR 4/2	253	10YR 5/6
164	2.5YR 6/10	209	7.5YR 3/2	254	10YR 8/8
165	2.5YR 5/10	210	7.5YR 9/4	255	10YR 7/8
166	2.5YR 7/12	211	7.5YR 8/4	256	10YR 6/8
167	2.5YR 6/12	212	7.5YR 7/4	257	10YR 5/8
168	2.5YR 6/14	213	7.5YR 6/4	258	10YR 8/10
169	5YR 9/1	214	7.5YR 5/4	259	10YR 7/10
170	5YR 8/1	215	7.5YR 4/4	260	10YR 6/10
171	5YR 7/1	216	7.5YR 8/6	261	10YR 7/12
172	5YR 6/1	217	7.5YR 7/6		
173	5YR 5/1	218	7.5YR 6/6	Index	YELLOW
174	5YR 4/1	219	7.5YR 5/6	262	2.5Y 9/2
175	5YR 3/1	220	7.5YR 4/6	263	2.5Y 8.5/2
176	5YR 2.5/1	221	7.5YR 8/8	264	2.5Y 8/2
177	5YR 9/2	222	7.5YR 7/8	265	2.5Y 7/2
178	5YR 8/2	223	7.5YR 6/8	266	2.5Y 6/2

Munsell Colour Nomenclature					
Index	YELLOW	Index	YELLOW	Index	YELLOW
267	2.5Y 5/2	312	5Y 7/4	357	7.5Y 7/8
268	2.5Y 4/2	313	5Y 6/4	358	7.5Y 6/8
269	2.5Y 3/2	314	5Y 5/4	359	7.5Y 9/10
270	2.5Y 9/4	315	5Y 4/4	360	7.5Y 8.5/10
271	2.5Y 8.5/4	316	5Y 9/6	361	7.5Y 8/10
272	2.5Y 8/4	317	5Y 8.5/6	362	7.5Y 7/10
273	2.5Y 7/4	318	5Y 8/6	363	7.5Y 8.5/12
274	2.5Y 6/4	319	5Y 7/6	364	7.5Y 8/12
275	2.5Y 5/4	320	5Y 6/6	365	10Y 9/1
276	2.5Y 4/4	321	5Y 5/6	366	10Y 8.5/1
277	2.5Y 9/6	322	5Y 9/8	367	10Y 8/1
278	2.5Y 8.5/6	323	5Y 8.5/8	368	10Y 7/1
279	2.5Y 8/6	324	5Y 8/8	369	10Y 6/1
280	2.5Y 7/6	325	5Y 7/8	370	10Y 5/1
281	2.5Y 6/6	326	5Y 6/8	371	10Y 4/1
282	2.5Y 5/6	327	5Y 8.5/10	372	10Y 3/1
283	2.5Y 8.5/8	328	5Y 8/10	373	10Y 2.5/1
284	2.5Y 8/8	329	5Y 7/10	374	10Y 9/2
285	2.5Y 7/8	330	5Y 8.5/12	375	10Y 8.5/2
286	2.5Y 6/8	331	5Y 8/12	376	10Y 8/2
287	2.5Y 8.5/10	332	5Y 7/12	377	10Y 7/2
288	2.5Y 8/10	333	7.5Y 9/2	378	10Y 6/2
289	2.5Y 7/10	334	7.5Y 8.5/2	379	10Y 5/2
290	2.5Y 8/12	335	7.5Y 8/2	380	10Y 4/2
291	2.5Y 7/12	336	7.5Y 7/2	381	10Y 3/2
292	5Y 9/1	337	7.5Y 6/2	382	10Y 9/4
293	5Y 8.5/1	338	7.5Y 5/2	383	10Y 8.5/4
294	5Y 8/1	339	7.5Y 4/2	384	10Y 8/4
295	5Y 7/1	340	7.5Y 3/2	385	10Y 7/4
296	5Y 6/1	341	7.5Y 9/4	386	10Y 6/4
297	5Y 5/1	342	7.5Y 8.5/4	387	10Y 5/4
298	5Y 4/1	343	7.5Y 8/4	388	10Y 4/4
299	5Y 3/1	344	7.5Y 7/4	389	10Y 9/6
300	5Y 2.5/1	345	7.5Y 6/4	390	10Y 8.5/6
301	5Y 9/2	346	7.5Y 5/4	391	10Y 8/6
302	5Y 8.5/2	347	7.5Y 4/4	392	10Y 7/6
303	5Y 8/2	348	7.5Y 9/6	393	10Y 6/6
304	5Y 7/2	349	7.5Y 8.5/6	394	10Y 5/6
305	5Y 6/2	350	7.5Y 8/6	395	10Y 9/8
306	5Y 5/2	351	7.5Y 7/6	396	10Y 8.5/8
307	5Y 4/2	352	7.5Y 6/6	397	10Y 8/8
308	5Y 3/2	353	7.5Y 5/6	398	10Y 7/8
309	5Y 9/4	354	7.5Y 9/8	399	10Y 6/8
310	5Y 8.5/4	355	7.5Y 8.5/8	400	10Y 9/10
311	5Y 8/4	356	7.5Y 8/8	401	10Y 8.5/10

Munsell Colour Nomenclature

Index	YELLOW
402	10Y 8/10
403	10Y 7/10
404	10Y 8/12

Index	GREEN YELLOW
405	2.5GY 9/2
406	2.5GY 8.5/2
407	2.5GY 8/2
408	2.5GY 7/2
409	2.5GY 6/2
410	2.5GY 5/2
411	2.5GY 4/2
412	2.5GY 3/2
413	2.5GY 9/4
414	2.5GY 8.5/4
415	2.5GY 8/4
416	2.5GY 7/4
417	2.5GY 6/4
418	2.5GY 5/4
419	2.5GY 4/4
420	2.5GY 9/6
421	2.5GY 8.5/6
422	2.5GY 8/6
423	2.5GY 7/6
424	2.5GY 6/6
425	2.5GY 5/6
426	2.5GY 9/8
427	2.5GY 8.5/8
428	2.5GY 8/8
429	2.5GY 7/8
430	2.5GY 6/8
431	2.5GY 8.5/10
432	2.5GY 8/10
433	2.5GY 7/10
434	5GY 9/1
435	5GY 8.5/1
436	5GY 8/1
437	5GY 7/1
438	5GY 6/1
439	5GY 5/1
440	5GY 4/1
441	5GY 3/1
442	5GY 2.5/1
443	5GY 9/2
444	5GY 8.5/2

Index	GREEN YELLOW
445	5GY 8/2
446	5GY 7/2
447	5GY 6/2
448	5GY 5/2
449	5GY 4/2
450	5GY 3/2
451	5GY 9/4
452	5GY 8.5/4
453	5GY 8/4
454	5GY 7/4
455	5GY 6/4
456	5GY 5/4
457	5GY 4/4
458	5GY 9/6
459	5GY 8.5/6
460	5GY 8/6
461	5GY 7/6
462	5GY 6/6
463	5GY 5/6
464	5GY 8.5/8
465	5GY 8/8
466	5GY 7/8
467	5GY 6/8
468	5GY 5/8
469	5GY 8.5/10
470	5GY 8/10
471	7.5GY 9/2
472	7.5GY 8.5/2
473	7.5GY 8/2
474	7.5GY 7/2
475	7.5GY 6/2
476	7.5GY 5/2
477	7.5GY 4/2
478	7.5GY 3/2
479	7.5GY 9/4
480	7.5GY 8.5/4
481	7.5GY 8/4
482	7.5GY 7/4
483	7.5GY 6/4
484	7.5GY 5/4
485	7.5GY 4/4
486	7.5GY 8.5/6
487	7.5GY 8/6
488	7.5GY 7/6
489	7.5GY 6/6

Index	GREEN YELLOW
490	7.5GY 7/8
491	7.5GY 6/8
492	7.5GY 5/8
493	7.5GY 5/6
494	7.5GY 8/8
495	7.5GY 7/10
496	7.5GY 6/10
497	10GY 9/1
498	10GY 8.5/1
499	10GY 8/1
500	10GY 7/1
501	10GY 6/1
502	10GY 5/1
503	10GY 4/1
504	10GY 3/1
505	10GY 2.5/1
506	10GY 9/2
507	10GY 8.5/2
508	10GY 8/2
509	10GY 7/2
510	10GY 6/2
511	10GY 5/2
512	10GY 4/2
513	10GY 3/2
514	10GY 2.5/2
515	10GY 9/4
516	10GY 8.5/4
517	10GY 8/4
518	10GY 7/4
519	10GY 6/4
520	10GY 5/4
521	10GY 4/4
522	10GY 8.5/6
523	10GY 8/6
524	10GY 7/6
525	10GY 6/6
526	10GY 5/6
527	10GY 7/8
528	10GY 6/8
529	10GY 5/8
530	10GY 7/10
531	10GY 6/10

Index	GREEN
532	2.5G 9/2

Munsell Colour Nomenclature					
Index	GREEN	Index	GREEN	Index	GREEN
533	2.5G 8/2	578	5G 4/4	623	10G 2.5/1
534	2.5G 7/2	579	5G 3/4	624	10G 9/2
535	2.5G 6/2	580	5G 8/6	625	10G 8/2
536	2.5G 5/2	581	5G 7/6	626	10G 7/2
537	2.5G 4/2	582	5G 6/6	627	10G 6/2
538	2.5G 3/2	583	5G 5/6	628	10G 5/2
539	2.5G 2.5/2	584	5G 4/6	629	10G 4/2
540	2.5G 8/4	585	5G 7/8	630	10G 3/2
541	2.5G 7/4	586	5G 6/8	631	10G 2.5/2
542	2.5G 6/4	587	5G 5/8	632	10G 8/4
543	2.5G 5/4	588	5G 4/8	633	10G 7/4
544	2.5G 4/4	589	5G 7/10	634	10G 6/4
545	2.5G 3/4	590	5G 6/10	635	10G 5/4
546	2.5G 8/6	591	7.5G 9/2	636	10G 4/4
547	2.5G 7/6	592	7.5G 8/2	637	10G 3/4
548	2.5G 6/6	593	7.5G 7/2	638	10G 7/6
549	2.5G 5/6	594	7.5G 6/2	639	10G 6/6
550	2.5G 4/6	595	7.5G 5/2	640	10G 5/6
551	2.5G 7/8	596	7.5G 4/2	641	10G 4/6
552	2.5G 6/8	597	7.5G 3/2	642	10G 7/8
553	2.5G 5/8	598	7.5G 2.5/2	643	10G 6/8
554	2.5G 7/10	599	7.5G 8/4	644	10G 5/8
555	2.5G 6/10	600	7.5G 7/4	645	10G 4/8
556	2.5G 5/10	601	7.5G 6/4	646	10G 6/10
557	2.5G 6/12	602	7.5G 5/4		
558	5G 9/1	603	7.5G 4/4	Index	BLUE GREEN
559	5G 8/1	604	7.5G 3/4	647	2.5BG 9/2
560	5G 7/1	605	7.5G 8/6	648	2.5BG 8/2
561	5G 6/1	606	7.5G 7/6	649	2.5BG 7/2
562	5G 5/1	607	7.5G 6/6	650	2.5BG 6/2
563	5G 4/1	608	7.5G 5/6	651	2.5BG 5/2
564	5G 3/1	609	7.5G 4/6	652	2.5BG 4/2
565	5G 2.5/1	610	7.5G 7/8	653	2.5BG 3/2
566	5G 9/2	611	7.5G 6/8	654	2.5BG 2.5/2
567	5G 8/2	612	7.5G 5/8	655	2.5BG 8/4
568	5G 7/2	613	7.5G 4/8	656	2.5BG 7/4
569	5G 6/2	614	7.5G 7/10	657	2.5BG 6/4
570	5G 5/2	615	7.5G 6/10	658	2.5BG 5/4
571	5G 4/2	616	10G 9/1	659	2.5BG 4/4
572	5G 3/2	617	10G 8/1	660	2.5BG 3/4
573	5G 2.5/2	618	10G 7/1	661	2.5BG 7/6
574	5G 8/4	619	10G 6/1	662	2.5BG 6/6
575	5G 7/4	620	10G 5/1	663	2.5BG 5/6
576	5G 6/4	621	10G 4/1	664	2.5BG 4/6
577	5G 5/4	622	10G 3/1	665	2.5BG 3/6

Munsell Colour Nomenclature					
Index	BLUE GREEN	Index	BLUE GREEN	Index	BLUE
666	2.5BG 7/8	711	7.5BG 6/4	754	2.5B 8/2
667	2.5BG 6/8	712	7.5BG 5/4	755	2.5B 7/2
668	2.5BG 5/8	713	7.5BG 4/4	756	2.5B 6/2
669	2.5BG 4/8	714	7.5BG 3/4	757	2.5B 5/2
670	2.5BG 6/10	715	7.5BG 7/6	758	2.5B 4/2
671	5BG 9/1	716	7.5BG 6/6	759	2.5B 3/2
672	5BG 8/1	717	7.5BG 5/6	760	2.5B 2.5/2
673	5BG 7/1	718	7.5BG 4/6	761	2.5B 8/4
674	5BG 6/1	719	7.5BG 7/8	762	2.5B 7/4
675	5BG 5/1	720	7.5BG 6/8	763	2.5B 6/4
676	5BG 4/1	721	7.5BG 5/8	764	2.5B 5/4
677	5BG 3/1	722	7.5BG 4/8	765	2.5B 4/4
678	5BG 2.5/1	723	10BG 9/1	766	2.5B 3/4
679	5BG 9/2	724	10BG 8/1	767	2.5B 7/6
680	5BG 8/2	725	10BG 7/1	768	2.5B 6/6
681	5BG 7/2	726	10BG 6/1	769	2.5B 5/6
682	5BG 6/2	727	10BG 5/1	770	2.5B 4/6
683	5BG 5/2	728	10BG 4/1	771	2.5B 7/8
684	5BG 4/2	729	10BG 3/1	772	2.5B 6/8
685	5BG 3/2	730	10BG 2.5/1	773	2.5B 5/8
686	5BG 2.5/2	731	10BG 9/2	774	2.5B 4/8
687	5BG 8/4	732	10BG 8/2	775	5B 9/1
688	5BG 7/4	733	10BG 7/2	776	5B 8/1
689	5BG 6/4	734	10BG 6/2	777	5B 7/1
690	5BG 5/4	735	10BG 5/2	778	5B 6/1
691	5BG 4/4	736	10BG 4/2	779	5B 5/1
692	5BG 3/4	737	10BG 3/2	780	5B 4/1
693	5BG 7/6	738	10BG 2.5/2	781	5B 3/1
694	5BG 6/6	739	10BG 8/4	782	5B 2.5/1
695	5BG 5/6	740	10BG 7/4	783	5B 9/2
696	5BG 4/6	741	10BG 6/4	784	5B 8/2
697	5BG 7/8	742	10BG 5/4	785	5B 7/2
698	5BG 6/8	743	10BG 4/4	786	5B 6/2
699	5BG 5/8	744	10BG 3/4	787	5B 5/2
700	5BG 4/8	745	10BG 7/6	788	5B 4/2
701	7.5BG 9/2	746	10BG 6/6	789	5B 3/2
702	7.5BG 8/2	747	10BG 5/6	790	5B 2.5/2
703	7.5BG 7/2	748	10BG 4/6	791	5B 8/4
704	7.5BG 6/2	749	10BG 7/8	792	5B 7/4
705	7.5BG 5/2	750	10BG 6/8	793	5B 6/4
706	7.5BG 4/2	751	10BG 5/8	794	5B 5/4
707	7.5BG 3/2	752	10BG 4/8	795	5B 4/4
708	7.5BG 2.5/2	753	2.5B 9/2	796	5B 3/4
709	7.5BG 8/4			797	5B 7/6
710	7.5BG 7/4			798	5B 6/6

Munsell Colour Nomenclature

Index	BLUE	Index	BLUE	Index	PURPLE BLUE
799	5B 5/6	844	10B 4/2	887	2.5PB 6/8
800	5B 4/6	845	10B 3/2	888	2.5PB 5/8
801	5B 3/6	846	10B 2.5/2	889	2.5PB 4/8
802	5B 7/8	847	10B 8/4	890	2.5PB 6/10
803	5B 6/8	848	10B 7/4	891	2.5PB 5/10
804	5B 5/8	849	10B 6/4	892	2.5PB 4/10
805	5B 4/8	850	10B 5/4	893	5PB 9/1
806	7.5B 9/2	851	10B 4/4	894	5PB 8/1
807	7.5B 8/2	852	10B 3/4	895	5PB 7/1
808	7.5B 7/2	853	10B 2.5/4	896	5PB 6/1
809	7.5B 6/2	854	10B 7/6	897	5PB 5/1
810	7.5B 5/2	855	10B 6/6	898	5PB 4/1
811	7.5B 4/2	856	10B 5/6	899	5PB 3/1
812	7.5B 3/2	857	10B 4/6	900	5PB 2.5/1
813	7.5B 2.5/2	858	10B 3/6	901	5PB 9/2
814	7.5B 8/4	859	10B 7/8	902	5PB 8/2
815	7.5B 7/4	860	10B 6/8	903	5PB 7/2
816	7.5B 6/4	861	10B 5/8	904	5PB 6/2
817	7.5B 5/4	862	10B 4/8	905	5PB 5/2
818	7.5B 4/4	863	10B 6/10	906	5PB 4/2
819	7.5B 3/4	864	10B 5/10	907	5PB 3/2
820	7.5B 7/6			908	5PB 2.5/2
821	7.5B 6/6	Index	PURPLE BLUE	909	5PB 8/4
822	7.5B 5/6	865	2.5PB 9/2	910	5PB 7/4
823	7.5B 4/6	866	2.5PB 8/2	911	5PB 6/4
824	7.5B 3/6	867	2.5PB 7/2	912	5PB 5/4
825	7.5B 7/8	868	2.5PB 6/2	913	5PB 4/4
826	7.5B 6/8	869	2.5PB 5/2	914	5PB 3/4
827	7.5B 5/8	870	2.5PB 4/2	915	5PB 2.5/4
828	7.5B 4/8	871	2.5PB 3/2	916	5PB 8/6
829	7.5B 6/10	872	2.5PB 2.5/2	917	5PB 7/6
830	7.5B 5/10	873	2.5PB 8/4	918	5PB 6/6
831	10B 9/1	874	2.5PB 7/4	919	5PB 5/6
832	10B 8/1	875	2.5PB 6/4	920	5PB 4/6
833	10B 7/1	876	2.5PB 5/4	921	5PB 3/6
834	10B 6/1	877	2.5PB 4/4	922	5PB 7/8
835	10B 5/1	878	2.5PB 3/4	923	5PB 6/8
836	10B 4/1	879	2.5PB 2.5/4	924	5PB 5/8
837	10B 3/1	880	2.5PB 8/6	925	5PB 4/8
838	10B 2.5/1	881	2.5PB 7/6	926	5PB 3/8
839	10B 9/2	882	2.5PB 6/6	927	5PB 6/10
840	10B 8/2	883	2.5PB 5/6	928	5PB 5/10
841	10B 7/2	884	2.5PB 4/6	929	5PB 4/10
842	10B 6/2	885	2.5PB 3/6	930	5PB 5/12
843	10B 5/2	886	2.5PB 7/8	931	5PB 4/12

Munsell Colour Nomenclature

Index	PURPLE BLUE	Index	PURPLE BLUE	Index	PURPLE
932	7.5PB 9/2	977	10PB 5/2	1020	2.5P 4/6
933	7.5PB 8/2	978	10PB 4/2	1021	2.5P 3/6
934	7.5PB 7/2	979	10PB 3/2	1022	2.5P 7/8
935	7.5PB 6/2	980	10PB 2.5/2	1023	2.5P 6/8
936	7.5PB 5/2	981	10PB 8/4	1024	2.5P 5/8
937	7.5PB 4/2	982	10PB 7/4	1025	2.5P 4/8
938	7.5PB 3/2	983	10PB 6/4	1026	2.5P 3/8
939	7.5PB 2.5/2	984	10PB 5/4	1027	2.5P 6/10
940	7.5PB 8/4	985	10PB 4/4	1028	2.5P 5/10
941	7.5PB 7/4	986	10PB 3/4	1029	2.5P 4/10
942	7.5PB 6/4	987	10PB 2.5/4	1030	5P 9/1
943	7.5PB 5/4	988	10PB 7/6	1031	5P 8/1
944	7.5PB 4/4	989	10PB 6/6	1032	5P 7/1
945	7.5PB 3/4	990	10PB 5/6	1033	5P 6/1
946	7.5PB 2.5/4	991	10PB 4/6	1034	5P 5/1
947	7.5PB 8/6	992	10PB 3/6	1035	5P 4/1
948	7.5PB 7/6	993	10PB 2.5/6	1036	5P 3/1
949	7.5PB 6/6	994	10PB 7/8	1037	5P 2.5/1
950	7.5PB 5/6	995	10PB 6/8	1038	5P 9/2
951	7.5PB 4/6	996	10PB 5/8	1039	5P 8/2
952	7.5PB 3/6	997	10PB 4/8	1040	5P 7/2
953	7.5PB 2.5/6	998	10PB 3/8	1041	5P 6/2
954	7.5PB 7/8	999	10PB 6/10	1042	5P 5/2
955	7.5PB 6/8	1000	10PB 5/10	1043	5P 4/2
956	7.5PB 5/8	1001	10PB 4/10	1044	5P 3/2
957	7.5PB 4/8			1045	5P 2.5/2
958	7.5PB 3/8	Index	PURPLE	1046	5P 8/4
959	7.5PB 6/10	1002	2.5P 9/2	1047	5P 7/4
960	7.5PB 5/10	1003	2.5P 8/2	1048	5P 6/4
961	7.5PB 4/10	1004	2.5P 7/2	1049	5P 5/4
962	7.5PB 3/10	1005	2.5P 6/2	1050	5P 4/4
963	7.5PB 5/12	1006	2.5P 5/2	1051	5P 3/4
964	7.5PB 4/12	1007	2.5P 4/2	1052	5P 2.5/4
965	10PB 9/1	1008	2.5P 3/2	1053	5P 7/6
966	10PB 8/1	1009	2.5P 2.5/2	1054	5P 6/6
967	10PB 7/1	1010	2.5P 8/4	1055	5P 5/6
968	10PB 6/1	1011	2.5P 7/4	1056	5P 4/6
969	10PB 5/1	1012	2.5P 6/4	1057	5P 3/6
970	10PB 4/1	1013	2.5P 5/4	1058	5P 2.5/6
971	10PB 3/1	1014	2.5P 4/4	1059	5P 7/8
972	10PB 2.5/1	1015	2.5P 3/4	1060	5P 6/8
973	10PB 9/2	1016	2.5P 2.5/4	1061	5P 5/8
974	10PB 8/2	1017	2.5P 7/6	1062	5P 4/8
975	10PB 7/2	1018	2.5P 6/6	1063	5P 3/8
976	10PB 6/2	1019	2.5P 5/6	1064	5P 5/10

Munsell Colour Nomenclature

Index	PURPLE	Index	PURPLE	Index	RED PURPLE
1065	5P 4/10	1110	10P 2.5/2	1154	2.5RP 7/8
1066	7.5P 9/2	1111	10P 8/4	1155	2.5RP 6/8
1067	7.5P 8/2	1112	10P 7/4	1156	2.5RP 5/8
1068	7.5P 7/2	1113	10P 6/4	1157	2.5RP 4/8
1069	7.5P 6/2	1114	10P 5/4	1158	2.5RP 7/10
1070	7.5P 5/2	1115	10P 4/4	1159	2.5RP 6/10
1071	7.5P 4/2	1116	10P 3/4	1160	2.5RP 5/10
1072	7.5P 3/2	1117	10P 2.5/4	1161	2.5RP 4/10
1073	7.5P 2.5/2	1118	10P 8/6	1162	2.5RP 6/12
1074	7.5P 8/4	1119	10P 7/6	1163	2.5RP 5/12
1075	7.5P 7/4	1120	10P 6/6	1164	5RP 9/1
1076	7.5P 6/4	1121	10P 5/6	1165	5RP 8/1
1077	7.5P 5/4	1122	10P 4/6	1166	5RP 7/1
1078	7.5P 4/4	1123	10P 3/6	1167	5RP 6/1
1079	7.5P 3/4	1124	10P 7/8	1168	5RP 5/1
1080	7.5P 2.5/4	1125	10P 6/8	1169	5RP 4/1
1081	7.5P 8/6	1126	10P 5/8	1170	5RP 3/1
1082	7.5P 7/6	1127	10P 4/8	1171	5RP 2.5/1
1083	7.5P 6/6	1128	10P 3/8	1172	5RP 9/2
1084	7.5P 5/6	1129	10P 6/10	1173	5RP 8/2
1085	7.5P 4/6	1130	10P 5/10	1174	5RP 7/2
1086	7.5P 3/6	1131	10P 4/10	1175	5RP 6/2
1087	7.5P 7/8	1132	10P 5/12	1176	5RP 5/2
1088	7.5P 6/8			1177	5RP 4/2
1089	7.5P 5/8	Index	RED PURPLE	1178	5RP 3/2
1090	7.5P 4/8	1134	2.5RP 8/2	1179	5RP 2.5/2
1091	7.5P 3/8	1135	2.5RP 7/2	1180	5RP 8/4
1092	7.5P 6/10	1136	2.5RP 6/2	1181	5RP 7/4
1093	7.5P 5/10	1137	2.5RP 5/2	1182	5RP 6/4
1094	7.5P 4/10	1138	2.5RP 4/2	1183	5RP 5/4
1095	10P 9/1	1139	2.5RP 3/2	1184	5RP 4/4
1096	10P 8/1	1140	2.5RP 2.5/2	1185	5RP 3/4
1097	10P 7/1	1141	2.5RP 8/4	1186	5RP 2.5/4
1098	10P 6/1	1142	2.5RP 7/4	1187	5RP 8/6
1099	10P 5/1	1143	2.5RP 6/4	1188	5RP 7/6
1100	10P 4/1	1144	2.5RP 5/4	1189	5RP 6/6
1101	10P 3/1	1145	2.5RP 4/4	1190	5RP 5/6
1102	10P 2.5/1	1146	2.5RP 3/4	1191	5RP 4/6
1103	10P 9/2	1147	2.5RP 2.5/4	1192	5RP 3/6
1104	10P 8/2	1148	2.5RP 8/6	1193	5RP 7/8
1105	10P 7/2	1149	2.5RP 7/6	1194	5RP 6/8
1106	10P 6/2	1150	2.5RP 6/6	1195	5RP 5/8
1107	10P 5/2	1151	2.5RP 5/6	1196	5RP 4/8
1108	10P 4/2	1152	2.5RP 4/6	1197	5RP 6/10
1109	10P 3/2	1153	2.5RP 3/6	1198	5RP 5/10

Munsell Colour Nomenclature					
Index	RED PURPLE	Index	RED PURPLE	Index	RED PURPLE
1199	5RP 4/10	1223	7.5RP 7/8	1247	10RP 2.5/2
1200	5RP 5/12	1224	7.5RP 6/8	1248	10RP 8/4
1201	5RP 4/12	1225	7.5RP 5/8	1249	10RP 7/4
1202	7.5RP 9/2	1226	7.5RP 4/8	1250	10RP 6/4
1203	7.5RP 8/2	1227	7.5RP 6/10	1251	10RP 5/4
1204	7.5RP 7/2	1228	7.5RP 5/10	1252	10RP 4/4
1205	7.5RP 6/2	1229	7.5RP 4/10	1253	10RP 3/4
1206	7.5RP 5/2	1230	7.5RP 5/12	1254	10RP 2.5/4
1207	7.5RP 4/2	1231	7.5RP 4/12	1255	10RP 8/6
1208	7.5RP 3/2	1232	10RP 9/1	1256	10RP 7/6
1209	7.5RP 2.5/2	1233	10RP 8/1	1257	10RP 6/6
1210	7.5RP 8/4	1234	10RP 7/1	1258	10RP 5/6
1211	7.5RP 7/4	1235	10RP 6/1	1259	10RP 4/6
1212	7.5RP 6/4	1236	10RP 5/1	1260	10RP 3/6
1213	7.5RP 5/4	1237	10RP 4/1	1261	10RP 7/8
1214	7.5RP 4/4	1238	10RP 3/1	1262	10RP 6/8
1215	7.5RP 3/4	1239	10RP 2.5/1	1263	10RP 5/8
1216	7.5RP 2.5/4	1240	10RP 9/2	1264	10RP 4/8
1217	7.5RP 8/6	1241	10RP 8/2	1265	10RP 6/10
1218	7.5RP 7/6	1242	10RP 7/2	1266	10RP 5/10
1219	7.5RP 6/6	1243	10RP 6/2	1267	10RP 4/10
1220	7.5RP 5/6	1244	10RP 5/2	1268	10RP 5/12
1221	7.5RP 4/6	1245	10RP 4/2	1269	10RP 4/12
1222	7.5RP 3/6	1246	10RP 3/2		

Bibliography

[1] IMS, HDRC VGA *Imager, Camera data and features*, Institute for Microelectronis Stuttgart, Stuttgart, August 2003.

[2] R. A. Normann and I. Perlman, "The effects of background illumination on the photoresponses of red and green cones," *Journal of Physiology*, vol. 286, pp. 491 – 507, Janauary 1979.

[3] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 2nd ed. Addison-Wesley Publishing Company, September 2002.

[4] L. W. Lai, C. H. Lai, and Y. C. King, "A novel logarithmic response CMOS image sensor with high output voltage swing and in-pixel fixed pattern noise reduction," in *IEEE Sensors Journal*, vol. 4, February 2004, pp. 122 – 126.

[5] B. E. Bayer, "Color imaging array," US Patent No. 3971065, July 1976.

[6] G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. John Wiley & Sons, Inc., New York, 1982.

[7] M. Čadík, M. Wimmer, L. Neumann, and A. Artusi, "Image attributes and quality for evaluation of tone mapping operators," in *Proceedings of Pacific Graphics 2006*, ser. 14th Pacific Conference on Computer Graphics and Applications. National Taiwan University Press, Oct 2006, pp. 35–44.

[8] D. Laschinski, "http://www.cs.huji.ac.il/ danix," Website, 2008, associate professor, School of Computer Science and Engineering, Hebrew University of Jerusalem, Israel.

[9] B. Wandel, *Foundations of Vision*. Sinauer Associates Inc, 1995.

[10] P. Debevec, "http://www.debevec.org/," Website, 2008, associate Director and Research Associate Professor, Graphics Research.

[11] Cypress Semiconductor, “www.cypress.com,” Website, August 2007.

[12] D. Joseph and S. Collins, “Modelling, calibration and correction of non-linear illumination dependent fixed pattern noise in logarithmic CMOS image sensors,” in *Proceedings of the 18th IEEE Instrumentation and Measurement Technology Conference*, vol. 2, May 2001, pp. 1296 – 301.

[13] ——, “Modeling, calibration, and correction of non-linear illumination dependent fixed pattern noise in logarithmic CMOS image sensors,” *IEEE Transactions on Instrumentation and Measurement*, vol. 51, no. 5, pp. 996 – 1001, 2002.

[14] M. Loose, K. Meier, and J. Schemmel, “CMOS image sensor with logarithmic response and self calibrating fixed pattern noise correction,” *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 3410, pp. 117 – 127, 1998.

[15] ——, “A self-calibrating single-chip CMOS camera with logarithmic response,” *IEEE Journal of Solid-State Circuits*, vol. 36, no. 4, pp. 586 – 596, 2001.

[16] Infotrends, “http://www.infotrends.com/,” Website, 2006.

[17] W. S. Boyle and G. Smith, “Charge coupled semiconductor devices,” Bell Laboratories, New Hill, NJ, Tech. Rep. 4, 1970.

[18] E. R. Fossum, “CMOS image sensors: Electronic camera-on-a-chip,” *IEEE Transactions on Electron Devices*, vol. 44, no. 10, pp. 1689 – 1698, 1997.

[19] M. Bigas, E. Cabruja, J. Forest, and J. Salvi, “Review of CMOS image sensors,” *Microelectronics Journal*, vol. 37, no. 5, pp. 433 – 451, 2006.

[20] K. Diefendorff, “CMOS image sensors challenge CCD’s,” *Microprocessor Report*, pp. 1 – 5, 22 June 1998.

[21] K. Yoon, C. Kim, B. Lee, and D. Lee, “Single-chip CMOS image sensor for mobile applications,” *Solid-State Circuits, IEEE Journal of*, vol. 37, no. 12, pp. 1839 – 1845, Dec. 2002.

[22] T. Hammadou, F. Boussaid, and M. Nilsson, “Low cost single chip CMOS camera for automotive application,” *Digest of Technical Papers - IEEE International Conference on Consumer Electronics*, pp. 44 – 45, 2002.

[23] B. Hosticka, W. Brockherde, A. Bussmann, T. Heimann, R. Jeremias, A. Kemma, C. Nitta, and O. Schrey, "CMOS imaging for automotive applications," *IEEE Transactions on Electron Devices*, vol. 50, no. 1, pp. 173 – 183, 2003.

[24] C. Ponchut, "Evaluation of an x-ray area detector based on a logarithmic CMOS camera," *IEEE Nuclear Science Symposium and Medical Imaging Conference*, vol. 2, pp. 9 – 241, 2000.

[25] A. Krymski, D. Van Blerkom, A. Andersson, N. Bock, B. Mansoorian, and E. R. Fossum, "High speed $500\,frames/s, 1024 \times 1024$ CMOS active pixel sensor," *IEEE Symposium on VLSI Circuits, Digest of Technical Papers*, pp. 137 – 138, 1999.

[26] A. I. Krymski, N. E. Bock, N. Tu, D. Van Blerkom, and E. R. Fossum, "A high-speed 240 $frames/s$, 4.1-Mpixel CMOS sensor," *IEEE Transactions on Electron Devices*, vol. 50, no. 1, pp. 130 – 135, 2003.

[27] A. I. Krymski and N. Tu, "A 9 $V/lux-s$, 5000 $frames/s$, 512×512 CMOS sensor," *IEEE Transactions on Electron Devices*, vol. 50, no. 1, pp. 136 – 143, 2003.

[28] X. Zhang and B. A. Wandell, "Color image fidelity metrics evaluated using image distortion maps," *Signal Processing*, vol. 70, no. 3, pp. 201 – 214, 1998.

[29] A. J. Blanksby and M. J. Loinaz, "Performance analysis of a color CMOS photogate image sensor," *IEEE Transactions on Electron Devices*, vol. 47, no. 1, pp. 55 – 64, 2000.

[30] S. G. Chamberlain and J. P. Y. Lee, "Novel wide dynamic range silicon photodetector and linear imaging array," *IEEE Journal of Solid-State Circuits*, vol. SC-19, no. 1, pp. 41 – 48, 1984.

[31] O. Yadid Pecht, R. Ginosar, and Y. Shacham Diamand, "A random access photodiode array for intelligent image capture," *Electron Devices, IEEE Transactions on*, vol. 38, no. 8, pp. 1772 – 1780, Aug. 1991.

[32] W. Yang, "Wide-dynamic-range, low-power photosensor array," *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, pp. 230 – 231, 1994.

[33] A. C. Luther, *Video Camera Technology*. Artech House Inc, 1998.

[34] G. Meynants, "A single-chip colour camera in CMOS technology," Ph.D. dissertation, IMEC, Belgium, 1998.

[35] O. Y. Pecht, "Wide dynamic range sensors," in *University of Negev Publications*. Israel: Ben-Gurion, 2000.

[36] J. Hynecek, "Low-noise and high-speed charge detection in high-resolution CCD's image sensors," *IEEE Transactions on Electron Devices*, vol. 44, no. 10, pp. 1679 – 1688, 1997.

[37] E. C. Fox, G. R. Allan, B. Li, D. Dattani, S. Kamasz, M. J. Kiik, Q. Tang, A. Pavlov, D. Dykaar, and S. G. Ingram, "A 640×480 CMOS image sensor for high speed imaging," *IEEE Workshop on Charged-Coupled Devices and Advanced Image Sensors*, 2001.

[38] D. X. D. Yang, A. E. Gamal, F. Boyd, and T. Hui, "A 640×512 CMOS image sensor with ultrawide dynamic range floating-point pixel-level ADC," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 12, pp. 1821 – 34, December 1999.

[39] G. Chapinal, S. A. Bota, M. Moreno, J. Palacin, and A. Herms, "A 128×128 CMOS image sensor with analog memory for synchronous image capture," *IEEE Sensors Journal*, vol. 2, no. 2, pp. 120 – 127, 2002.

[40] Y. Muramatsu, S. Kurosawa, M. Furumiya, H. Ohkubo, and Y. Nakashiba, "A signal-processing CMOS image sensor using a simple analog operation," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 1, pp. 101 – 106, Jan. 2003.

[41] X. C. Guo, "A time based asynchronous readout CMOS image sensor," Ph.D. dissertation, University of Florida, 2002.

[42] S. Decker, R. D. McGrath, K. Brehmer, and C. G. Sodini, "A 256×256 CMOS imaging array with wide dynamic range pixels and column-parallel digital output," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 12, pp. 2081 – 2091, 1998.

[43] C. Xu, W. Q. Zhang, W. H. Ki, and M. Chan, "A 1.0V - V_{DD} CMOS active-pixel sensor with complementary pixel architecture and pulsewidth modulation fabricated with a $0.25\mu m$ CMOS process," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 12, pp. 1853 – 1859, 2002.

[44] S. H. Lim and A. El Gamal, "Gain fixed pattern noise correction via optical flow," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 51, no. 4, pp. 779 – 786, 2004.

[45] M. Schanz, C. Nitta, A. Bubmann, B. J. Hostika, and R. K. Wertheimer, "A high dynamic range CMOS image sensor for automotive applications." *IEEE Journal on Solid State Circuits*, vol. 35, no. 7, pp. 932 – 938, July 2000.

[46] T. Lule, M. Wagner, M. Verhoeven, H. Keller, and M. Boehm, "100,000-pixel, 120 – dB imager in TFA technology," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 5, pp. 732 – 739, 2000.

[47] D. Stoppa, A. Simoni, L. Gonzo, M. Gottardi, and G.-F. Dalla Betta, "Novel CMOS image sensor with a 132-dB dynamic range," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 12, pp. 1846 – 1852, 2002.

[48] Q. Luo and J. G. Harris, "A time-based CMOS image sensor," *Proceedings - IEEE International Symposium on Circuits and Systems*, vol. 4, pp. 840 – 843, 2004.

[49] S. F. Chen, Y. J. Juang, S. Y. Huang, and Y. C. King, "Logarithmic CMOS image sensor through multi-resolution analog-to-digital conversion," in *International Symposium on VLSI Technology, Systems, and Applications, 2003.*, Department of Electrical Engineering, National Tsing Hua University, Hsinchu, Taiwan. IEEE, Piscataway, NJ, USA, 2003, pp. 227 – 230.

[50] Y. C. Chuang, S. F. Chen, S. Y. Huang, and Y. C. King, "Low-cost logarithmic CMOS image sensing by non-linear analog-to-digital conversion," *IEEE Transactions on Consumer Electronics*, vol. 51, no. 4, pp. 1212 – 1217, 2005.

[51] C. Mead, *Analog* VLSI *and Neural Systems*. Addison - Wesley, Reading MA, 1989.

[52] T. Delbruck and C. A. Mead, "Adaptive photoreceptor with wide dynamic range," *Proceedings - IEEE International Symposium on Circuits and Systems*, vol. 4, pp. 339 – 342, 1994.

[53] IMEC, *Document Guide - Fuga 15*RGB *Camera*, C-Cam Technologies, April 1998.

[54] A. Bermak, A. Bouzerdoum, and K. Eshraghian, "High fill-factor native logarithmic pixel: Simulation, design and layout optimization," *Proceedings - IEEE International Symposium on Circuits and Systems*, vol. 5, pp. 293 – 296, 2000.

[55] K. Spyros, D. Bart, S. Danny, A. Andre, U. Dirk, and B. Jan, "A logarithmic response CMOS image sensor with on-chip calibration," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 8, pp. 1146 – 52, August 2000.

[56] D. P. Foty, *MOSFET Modelling with SPICE: Principles and Practice*. Upper Saddle River, NJ: Prentice Hall, 1997.

[57] Y. Tsividis, *Operation and Modelling of the* MOS *Transistor*, 2nd ed. McGraw Hill, 1999.

[58] S. O. Otim, B. Choubey, D. Joseph, and S. Collins, "Characterization and simple fixed pattern noise correction in wide dynamic range "logarithmic" imagers," *IEEE Transactions on Instrumentation and Measurement*, vol. 56, no. 5, pp. 1910–1916, October 2007.

[59] O. Vietze and P. Seitz, "Active pixels for image sensing with programmable, high dynamic range," *AT Proceedings: Advanced Technologies Intelligent Vision*, pp. 15 – 18, 1995.

[60] N. Ricquier and B. Dierickx, "Active pixel CMOS image sensor with on-chip non-uniformity correction," in *IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors*, April 1995.

[61] B. Dierickx, D. Scheffer, G. Meynants, W. Ogiers, and J. Vlummens, "Random addressable active pixel image sensors," in *Proceedings of SPIE*, vol. 2950, October 1996, pp. 2 – 7.

[62] G. Marshall and S. Collins, "A high dynamic range front end for automatic image processing applications," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 3410, pp. 176 – 185, 1998.

[63] S. Collins, J. Ngole, and G. Marshall, "High gain trimmable logarithmic CMOS pixel," *Electronics Letters*, vol. 36, no. 21, pp. 1806 – 1807, 2000.

[64] K. Spyros, D. Bart, and S. Danny, "On-chip offset calibrated logarithmic response image sensor," in *IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors*, June 1999, pp. 68 – 71.

[65] D. Scheffer, B. Dierickx, and G. Meynants, "Random addressable 2048×2048 active pixel image sensor," *IEEE Transactions on Electron Devices*, vol. 44, no. 10, pp. 1716 – 20, October 1997.

[66] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec, *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann Publishers, December 2005.

[67] R. M. Boynton, *Human Colour Vision*. San Diego: University of California, 1979.

[68] G. D. Finlayson and S. D. Hordley, "Color constancy at a pixel," *Journal of the Optical Society of America A: Optics and Image Science, and Vision*, vol. 18, no. 2, pp. 253 – 264, 2001.

[69] G. Finlayson, S. Hordley, and P. Hubel, "Color by correlation: A simple, unifying framework for color constancy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1209 – 1221, 2001.

[70] F. Graham and X. Ruixia, "Illuminant and gamma comprehensive normalisation in log RGB space," *Pattern Recogn. Lett.*, vol. 24, no. 11, pp. 1679 – 1690, 2003.

[71] G. Jan Mark, R. van den Boomgaard, A. W. M. Smeulders, and G. Theo, "Color constancy from physical principles," *Pattern Recognition Letters*, vol. 24, no. 11, pp. 1653 – 1662, 2003.

[72] D. H. Brainard and B. A. Wandell, "Analysis of the retinex theory of color vision," *Journal of Optical Society of America*, pp. 208 – 218, 1992.

[73] G. Finlayson and G. Schaefer, "Hue that is invariant to brightness and gamma," in *British Machine Vision Conference*, 2001, pp. 303 – 312.

[74] L. Pradeep, G. Michael, E. Pecht, and H. B, *Influence of Temperature on Microelectronics and System Reliability*, 1st ed. CRC Press, 1997.

[75] D. Joseph, "Modelling and calibration of logarithmic CMOS image sensors," Ph.D. dissertation, Oxford University, Keble College, 30th September 2002.

[76] D. Joseph and S. Collins, "Temperature dependence of fixed pattern noise in logarithmic cmos image sensors," *Conference Record - IEEE Instrumentation and Measurement Technology Conference*, p. 4258059, 2007.

[77] B. Choubey, S. Aoyama, D. Joseph, S. Otim, and S. Collins, "An electronic calibration scheme for logarithmic CMOS pixels," *Proceedings - IEEE International Symposium on Circuits and Systems*, vol. 4, pp. 856 – 859, 2004.

[78] B. Choubey, S. Aoyoma, S. Otim, D. Joseph, and S. Collins, "An electronic calibration scheme for logarithmic CMOS pixels," *IEEE Sensors Journal*, vol. 6, no. 4, pp. 950 – 956, 2006.

[79] G. Storm, J. Hurwitz, D. Renshaw, K. Findlater, R. Henderson, and M. Purcell, "Combined linear-logarithmic CMOS image sensor," *Digest of Technical Papers - IEEE International Solid-State Circuits Conference*, vol. 47, pp. 116 – 117, 2003.

[80] S. O. Otim, D. Joseph, B. Choubey, and S. Collins, "Modelling of high dynamic range logarithmic CMOS image sensors," *IEEE Instrumentation and Measurement Technology Conference*, vol. 1, pp. 451 – 456, 2004.

[81] C. Enz, F. Krummenacher, and E. Vittoz, "An analytical MOS transistor model valid in all regions of operation and dedicated to low voltage and low current applications," *Analogue Integrated Circuits and Signal Processing*, pp. 83 – 114, July 1995.

[82] M. M. Levine, *Fundamentals of Sensation and Perception*. Oxford University Press, 2000.

[83] S. O. Otim, B. Choubey, D. Joseph, and S. Collins, "Simplified fixed pattern noise correction for logarithmic sensors," *IEEE Workshop on Charged-Coupled Devices and Advanced Image Sensors*, pp. 80–83, June 9th -11th 2005.

[84] K. Asim and C. Roy, *Modern Concepts of Colour and Appearance*. Enfield, NH, 2000.

[85] C. Reis, "That foveon digital photo camera: The technology behind it," *Advanced Imaging Magazine*, vol. 15, no. 11, pp. 28 – 30, 2000.

[86] R. F. Lyon and P. M. Hubel, "Eyeing the camera: Into the next century," *Proceedings of the Color Imaging Conference: Color Science, Systems, and Applications*, pp. 349 – 355, 2002.

[87] P. M. Hubel, J. Liu, and R. J. Guttosch, "Spatial frequency response of color image sensors: Bayer color filters and foveon $\times 3$," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 5301, pp. 402 – 407, 2004.

[88] D. Alleysson, S. Susstrunk, and J. Herault, "Color demosaicing by estimating luminance and opponent chromatic signals in the fourier domain," *Proceedings of the Color Imaging Conference: Color Science, Systems and Applications*, pp. 331 – 336, 2002.

[89] D. Alleysson, S. Susstrunk, and J. Marguier, "Influence of spectral sensitivity functions on color demosaicing," *Proceedings of the Color Imaging Conference: Color Science, Systems and Applications*, pp. 351 – 357, 2003.

[90] D. Travis, *Effective Color Displays. Theory and Practice*. Academic Press, 1991.

[91] X. Zhang and B. A. Wandell, "Spatial extension of CIELAB for digital color-image reproduction," *Journal of the Society for Information Display*, vol. 5, no. 1, pp. 61 – 63, 1997.

[92] S. S. Guan and M. R. Lou, "Investigation of parametric effects using small colour differences," *Color Research and Application*, vol. 24, no. 5, pp. 331 – 343, 1999.

[93] H. R. Kang, *Color technology for electronic imaging devices*, 1st ed. Society of Photo-Optical Instrumentation Engineers (SPIE), Mar 1997.

[94] M. Mahy, L. Van Eycken, and A. Oosterlinck, "Evaluation of uniform color spaces developed after the adoption of CIELAB and CIELUV," *Color Research and Application*, vol. 19, no. 2, p. 105, 1994.

[95] A. Abrardo, V. Cappellini, M. Cappellini, and A. Mecocci, "Art-works colour calibration by using the vasari scanner," *Proceedings of the Color Imaging Conference: Color Science, Systems, and Applications*, pp. 94 – 97, 1997.

[96] J. Y. Hardeberg, "Acquisition and reproduction of colour images: colorimetric and multispectral approaches," Ph.D. dissertation, Ecole Nationale Suprieure des Tlcommunications, Paris, France, 1999.

[97] Matt Munsell Colour Research Database, "http://spectral.joensuu.fi/databases/," University of Joensuu, Colour Research Website, February 2004.

[98] International Telecommunications Union (ITU-R), "Basic parameter values for the hdtv standard for the studio and for international programme exchange," ITU-R Recommendation BT.709, Geneva 20, Switzerland, 1990.

[99] A. K. Romney and T. Indow, "Munsell reflectance spectra represented in three dimensional euclidean space," *Color Research and Application*, vol. 28, no. 3, pp. 182 – 196, 2003.

[100] R. G. D'Andrade and A. K. Romney, "A quantitative model for transforming reflectance spectra into the munsell color space using cone sensitivity functions and opponent process weights." *Proceedings of the National Academy of Science United States of America*, vol. 100, no. 10, pp. 6281 – 6, May 2003.

[101] A. Romney, R. D'Andrade, and T. Indow, "The distribution of response spectra in the lateral geniculate nucleus compared with reflectance spectra of munsell color chips." *Proceedings of National Academy of Science USA*, vol. 102, no. 27, pp. 9720 – 5, 2005 Jul 5.

[102] L. Patrick, "Tone mapping for high dynamic range imaging," Ph.D. dissertation, University of Bristol, 2005.

[103] M. Kresimir, N. Laszlo, and P. Werner, "A survey of tone mapping techniques," Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, Tech. Rep. TR-186-2-97-12, April 1997.

[104] G. W. Larson, H. Rushmeier, and C. Piatko, "Visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 3, no. 4, pp. 291 – 306, 1997.

[105] M. S. Livingstone, *Vision and Art: The Biology of Seeing*. Harry N Abrams, Inc Publishers, 2002.

[106] P. Ledda, G. Ward, and A. Chalmers, "A wide field, high dynamic range, stereographic viewer," in *GRAPHITE '03: Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia*. New York, NY, USA: ACM Press, 2003, pp. 237 – 244.

[107] H. Seetzen, L. Whitehead, and G. Ward, "A high dynamic range display system using low and high resolution modulators," *Proceedings of the 2003 Society for Information Display Symposium*, 2003.

[108] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 760 – 768, 2004.

[109] P. Ledda, A. Chalmers, and H. Seetzen, "HDR displays: A validation against reality," *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, vol. 3, pp. 2777 – 2782, 2004.

[110] P. Ledda, A. Chalmers, T. Troscianko, and H. Seetzen, "Evaluation of tone mapping operators using a high dynamic range display," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 640 – 648, 2005.

[111] Dolby Digital (Brightside Technologies), "www.dolby.com/promo/hdr/," Website, August 2005.

[112] S. Oh, "High dynamic range image encoding for brightside display," Department of Electrical Engineering, Tech. Rep., March 2007.

[113] D. Kate, "A review of tone reproduction techniques," Department of Computer Science, University of Bristol, Tech. Rep. CSTR-02-005, November 2002.

[114] J. Tumblin and H. Rushmeier, "Tone reproduction for realistic images," in *IEEE Computer Graphics and Applications*, vol. 13(6), November 1993, pp. 42 – 46.

[115] G. Ward, "A contrast-based scalefactor for luminance display," *Graphics Gems IV*, pp. 415 – 421, 1994.

[116] J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg, "Model of visual adaptation for realistic image synthesis," *Proceedings of the ACM SIGGRAPH Conference on Computer Graphics*, pp. 249 – 258, 1996.

[117] N. P. Sumanta, T. Jack, Y. Hector, and D. P. Greenberg, "Time-dependent visual adaptation for fast realistic image display," in *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Company, 2000, pp. 47 – 54.

[118] C. Schlick, "Quantisation techniques for visualization of high dynamic range pictures," in *Photorealistic Rendering Techniques, Springer, Berlin*, 1995, pp. 7 – 20.

[119] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman, "Spatially non-uniform scaling functions for high contrast images," *Proceedings - Graphics Interface*, pp. 245 – 253, 1993.

[120] D. J. Jobson, Z. Rahman, and G. A. Woodell, "Multiscale retinex for bridging the gap between color images and the human observation of scenes," *IEEE Transactions on Image Processing*, vol. 6, no. 7, pp. 965 – 976, 1997.

[121] T. Jack, J. K. Hodgins, and B. K. Guenter, "Two methods for display of high contrast images," *ACM Trans. Graph.*, vol. 18, no. 1, pp. 56 – 94, 1999.

[122] C. Jonathan, T. Chris, H. Tim, and P. E. Debevec, "Real-time high dynamic range texture mapping," in *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*. London, UK: Springer-Verlag, 2001, pp. 313 – 320.

[123] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 249 – 256, 2002.

[124] S. O. Otim and S. Collins, "Display of HDR images from logarithmic CMOS sensors," *ICIS '06: 30th International Congress of Imaging Science - Final Program and Proceedings*, pp. 164 – 167, 2006.

[125] D. A. Forsyth, "A novel algorithm for color constancy," *International Journal of Computer Vision*, vol. 5, no. 1, pp. 5 – 36, 1990.

[126] G. D. Finlayson and S. Susstrunk, "Color ratios and chromatic adaptation," *European Conference on Colour in Graphics, Imaging and Vision (CGIV)*, pp. 7 – 10, 2002.

[127] P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs,” *Proceedings of the ACM SIGGRAPH Conference on Computer Graphics*, pp. 369 – 378, 1997.

[128] M. Stokes, M. D. Fairchild, and R. S. Berns, “Precision requirements for digital color reproduction,” *ACM Transactions on Graphics*, vol. 11, no. 4, pp. 406 – 422, 1992.

[129] G. Cairns, C. Dachs, M. Brownlow, Y. Kubota, H. Washio, and M. Hijigawa, “Multi-format digital display with content driven display format,” Sharp Laboratories of Europe and Sharp Corporation LCD Development Group, Tech. Rep., 2001.