Illuminant and Device Invariance Using Histogram Equalisation

G. D. Finlayson[1]*, S. D. Hordley*[1]*, G. Schaefer*[2] *and G. Y. Tian*[3]
[1] *School of Computing Sciences, University of East Anglia, Norwich, UK*
[2] *The Faculty of Computing and Technology, The Nottingham Trent University, Nottingham, UK*
[3] *The University of Huddersfield, Queensgate, Huddersfield, UK*

Abstract

In this paper we propose a new device and illumination invariant image representation based on an existing greyscale image enhancement technique: histogram equalisation. Our method is based on the premise that the rank ordering of sensor responses is preserved across a change in imaging conditions (lighting or device). We set out the theoretical conditions under which this premise is true and we present empirical evidence which demonstrates that rank ordering is maintained in practice for a wide range of illuminants and imaging devices. We then show how we can exploit this rank invariance using histogram equalisation to derive an invariant image representation. Device and illuminant invariance are important in many imaging applications and in this paper we demonstrate the practical benefits of our new method in one such situation: the problem of image retrieval. We show that using the new invariant image representation to index into a database of images taken with a variety of devices under different lights provides very good indexing performance across all imaging conditions.

1. Introduction

Colour (*RGB*) images provide information which can potentially help in solving a wide range of imaging problems. For example it has been demonstrated [15, 12, 1] that characterising an image by the distribution of its colours (*RGB*s) is an effective way to locate images with similar content from amongst a diverse database of images. Colour has also been found to be useful for tasks such as image segmentation and object tracking. Using colour in this way makes the implicit assumption that the colours recorded by devices are an inherent property of the imaged objects and thus a reliable cue to their identity. In fact, *RGB*s recorded by any imaging device are a measure of the light reflected from the surface of an object and so depend in equal measure on both the surface properties of the imaged objects and the light under which they are illuminanted. So, an object which is lit by an illuminant which is itself reddish will be recorded by a camera as more red than will the same object lit under a more bluish illuminant. In addition image colour also depends on the properties of the recording device. Importantly, different imaging devices have different sensors which implies that an object which produces a given *RGB* response in one camera might well produce a quite different response in a different device.

To overcome these problems researchers have sought modified image representations such that one or more of these dependencies are removed. Until now research has concentrated on accounting for only illumination dependence and proposed algorithms can be classified as either *colour invariant* [10, 9, 5, 6, 14] or *colour constancy* algorithms [11, 8]. Colour constancy algorithms, which seek to recover a true estimate of object colour can be practically useful given simple imaging conditions and a well characterised imaging device [11]. There do exist algorithms which operate without knowledge of the imaging device but it has been shown [8] they do not provide a good enough estimate of object colour to be practically useful. Colour invariant methods seek new features, algebraic transforms of the original image which are illuminant invariant. While often simpler than colour constancy methods they also suffer from the fact they do not yet provide sufficient invariance to be practically useful [3]. Furthermore, neither approach addresses the issue of device invariance.

In this paper we present a new invariant image representation which addresses some of the weaknesses of previous work. In particular the method we present is both illumination independent and (in many cases) also device independent. Our method is based on the observation that while a change in illumination or device leads to significant changes in the recorded *RGB*s, the *rank orderings* of the responses of a given sensor are largely preserved. In fact, we show in this paper (§3) that under certain simplifying assumptions invariance of rank ordering follows directly from the image formation equation. In addition we present an empirical study (§4) which reveals that the preservation of rank ordering holds in practice both across a wide range of illuminants and a variety of imaging devices. Thus, an image representation which is based on rank ordering of recorded *RGB*s rather than on the *RGB*s themselves offers the possibility of accounting for both illumination and device dependence at the same time.

To derive an image representation which depends only on rank orderings we borrow a tool which has long been used by the image processing community for a quite different purpose. The technique is histogram equalisation which is typically applied to grey-scale images to produce an "enhanced" image. But in a departure from traditional image processing practice we apply the procedure not to a grey-scale image, but rather to each of the *R*, *G*, and *B* channels of a colour image independently of one another. We show that provided two images differ in such a way as to preserve the rank ordering of pixel values in each of the three channels then an application of histogram equalisation to each of the channels of the two images results in a pair of equivalent images. Thus provided a change in illuminant or device preserves rank ordering of pixel responses the application of histogram equalisation will provide us with an invariant representation of a scene which might subsequently be of use in a range of imaging applications.

Of course the reader may be surprised that we propose something so simple: histogram equalisation is a common tool. Paradoxically however, histogram equalising *R*, *G*, and *B* channels of an image is generally discouraged because this results in unnatural pseudo-colours. In the context of many imaging applications however, such pseudo-colours suffice since we are not interested in viewing the image but rather in extracting useful information from it. We demonstrate that the method does indeed recover useful information by applying it to the problem of image retrieval (§5). We show that the method out performs all previous approaches providing very good indexing across devices and close to perfect indexing across a change in illumination.

2. Background

We begin by adopting a simple model of image formation in which a scene is illuminated by a single light characterised by its spectral power distribution which we denote $E(\lambda)$ and which specifies how much energy the source emits at each wavelength (λ) of the electromagnetic spectrum. A surface in the scene is characterised by its spectral reflectance function $S(\lambda)$ which defines what proportion of light incident upon it the surface reflects on a per-wavelength basis. We assume that the light reflected from the scene is sampled by a 2-d array of sensors at the imaging device (a camera, or the eye) and that each sensor is characterised by $R_k(\lambda)$, its spectral sensitivity function which specifies how much light energy it absorbs at each wavelength of the spectrum. We assume that light from each point in the scene is sampled by three different classes of sensor ($k = 1,2,3$) and the response of each sensor is given by:

$$p_k = \int_{\omega} E(\lambda)S(\lambda)R_k(\lambda)d\lambda, \quad k = 1,\ldots,3 \qquad (1)$$

where the integral is taken over the range of wavelengths ω: the range for which the sensor has non-zero sensitivity. Thus an imaging device's response to light from a point in the scene is given by the triplet (p_1, p_2, p_3) which we also refer to in this paper as R, G, and B or just *RGB*. An image is thus a collection of *RGB*s representing the device's response to light from a range positions in the scene.

Equation (1) makes clear the fact that a device response depends both on properties of the sensor (it depends on $R_k(\lambda)$ and also on the prevailing illumination (on $E(\lambda)$). That is, responses are both device and illumination dependent. It follows that if no account is taken of these dependencies, an *RGB* cannot correctly considered to be an intrinsic property of an object.

To deal with the problem of illumination dependence we could attempt to estimate the illumination ($E(\lambda)$) and subsequently correct recorded *RGB*s to render responses colour constant: i.e. stable across a change in illumination. In practice estimating the scene illuminant is non-trivial and Funt *et al* [8] demonstrated that no colour constancy algorithm is sufficiently accurate to make such an approach useful in practical applications. More recent work [11] has shown that for simple imaging conditions and given good device calibration the colour constancy approach can work but unfortunately well calibrated devices are often not available.

A different approach is to derive from the image data some new representation (feature) of the image which is invariant to illumination. Such approaches are classified as colour (or illuminant) invariant approaches and many invariant features have been proposed [9, 5, 10, 14]. It is clear from Equation (1), that the interaction between light, surface and sensor is complex and invariant representations are formulated on the basis of further simplifications to this equation. A common approach is to adopt the so called *diagonal* model of illumination change in which it is proposed that sensor responses under a pair of illuminants are related by a diagonal matrix transform:

$$\begin{pmatrix} R^c \\ G^c \\ B^c \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^o \\ G^o \\ B^o \end{pmatrix} \qquad (2)$$

where the superscripts o and c characterise the pair of illuminants. The model is widely used, and has been shown to be well justified under many conditions [7]. Adopting such a model one simple illuminant invariant representation of an image can be derived by applying the following transform:

$$R' = \frac{R}{R_{ave}},\ G' = \frac{G}{G_{ave}},\ B' = \frac{B}{B_{ave}} \tag{3}$$

where the triplet $(R_{ave}, G_{ave}, B_{ave})$ denotes the mean of all *RGB*s in an image. It is easy to show that this so called *Greyworld* representation of an image is illumination invariant provided that Equation (2) holds.

Many other illuminant invariant representations have been derived, in some cases [10] by adopting different models of image formation. All derived invariants however share two common failings: first it has been demonstrated that when applied to the practical problem of image retrieval [3] none of these invariants affords good enough performance across a change in illumination. Secondly, none of these approaches considers the issue of device invariance.

Variation across devices occurs because different devices have different spectral sensitivity functions (different R_k in Equation (1)) and also because the colours recorded by a device are often not linearly related to scene radiance as Equation (1) suggests, but rather are some non-linear transform of this:

$$p_k = f\left(\int_{\omega} E(\lambda)S(\lambda)R_k(\lambda)d\lambda\right),\quad k = 1,\ldots,3 \tag{4}$$

The transform $f()$ is deliberately applied to *RGB* values recorded by a device for a number of reasons. First, many captured images will eventually be displayed on a monitor. Importantly colours displayed on a screen are not a linear function of the RGBs sent to the monitor. Rather, there exists a power function relationship between the incoming voltage and the displayed intensity. This relationship is known as the gamma of the monitor, where gamma describes the exponent of the power function [13]. To compensate for this gamma function images are usually stored in a way that reverses the effect of this transformation: that is by applying a power function with exponent of $1/\gamma$, where γ describes the gamma of the monitor, to the image RGBs. Unfortunately monitor gammas are not unique and so images from two different devices will not necessarily have the same gamma correction applied. In addition to gamma correction other more general non-linear "tone curve" corrections are often applied to images so as to change image contrast with the intention of creating a visually more pleasing image. Such transformations are device, and quite often, image dependent. How then are we to deal with these problems? In the next section we provide a first step in answering this question by demonstrating that while *RGB*s themselves change across device and illumination, the relative ordering of these responses remains fixed.

3. Rank Invariance of Sensor Responses

Let us consider again the diagonal model of illumination change defined by Equation (2). We observe that one implication of this model is that the rank ordering of sensor responses is preserved under a change of illumination. To see this, consider the responses to a single sensor R, such that R_i^o represents the response to a surface i under an illuminant o. Under a second illuminant, which we denote c, the surface will have response R_i^c and the pair of sensor responses are related by:

$$R_i^c = \alpha R_i^o \tag{5}$$

Equation (5) is true for all surfaces (that is, $\forall\ i$). Now, consider a pair of surfaces, i and j, viewed under illuminant o and suppose that $R_i^o > R_j^o$, then it follows from Equation (5) that:

$$R_i^o > R_j^o \Rightarrow \alpha R_i^o > \alpha R_j^o \Rightarrow R_i^c > R_j^c\ \forall\ i,j,\ \forall\ \alpha > 0 \tag{6}$$

That is, the rank ordering of sensor responses within a given channel is invariant to a change in illumination.

Next, consider the more general model of image formation (Equation (4)) in which sensor responses are allowed to undergo a possibly non-linear transformation. Rank ordering is also preserved in this case for a certain class of functions $f()$. Specifically, rank ordering is preserved provided that $f()$ is a monotonic increasing function. Importantly many of the transformations such as gamma or tone-curve corrections which are applied to images, satisfy this condition of monotonicity and are thus rank invariant. For example power (gamma) function transformations are rank invariant since:

$$R_i > R_j \Rightarrow (R_i)^{\gamma} > (R_j)^{\gamma}\quad \forall\ \ \gamma > 0 \tag{7}$$

It makes sense that tone-curve corrections applied to images should also be monotonic (and thus rank invariant) since such corrections are essentially mappings from input pixel values to output values. If this mapping is not monotonic then it can happen that two quite different input pixel values are mapped to the same output value.

This analysis is not fully general: it is possible given what we have said so far, that the responses from two devices whose spectral sensitivities are quite different might have a different ordering of responses. We investigate this point further in the next section with an empirical analysis of data from a wide class of imaging devices.

3.1. Rank Invariance in Practice

To assess rank ordering invariance we conducted an experiment similar to that of Dannemiller [2] who investigated the invariance of human cone cells under a change in illumination. He found that to a very good approximation rank orderings were maintained. Here, we investigate both a change in illumination and device. To investigate the invariance of rank orderings of sensor responses for a single device under changing illumination we proceed as follows. Let R_k represent the spectral sensitivity of the kth sensor of the device we wish to investigate. Now suppose we calculate (according to Equation (1)) the responses of this sensor to a set of surface reflectance functions under a fixed illuminant $E^1(\lambda)$. We denote those responses by the vector P_k^1. Similarly we denote by P_k^2 the responses of the same sensor to the same surfaces viewed under a second illuminant $E^2(\lambda)$. Next we define a function *rank*() which takes a vector argument and returns a vector whose elements contain the rank of the corresponding element in the argument. Then, if sensor responses are invariant to the illuminants E^1 and E^2, the following relationship must hold:

$$rank\left(P_k^1\right) = rank\left(P_k^2\right) \tag{8}$$

In practice the relationship in Equation (8) will hold only approximately and we can assess how well the relationship holds using Spearman's Rank Correlation coefficient which is given by:

$$\rho = 1 - 6\sum_{j=1}^{N} \frac{d_j^2}{N_s(N_s^2-1)} \tag{9}$$

where d_j is the difference between the jth elements of $rank\left(P_k^1\right)$ and $rank\left(P_k^2\right)$ and N_s is the number of surfaces. This coefficient takes a value between -1 and 1: a coefficient of zero implies that Equation (8) holds not at all, while a value of one will be obtained when the relationship is exact. Invariance of rank ordering across devices can be assessed in a similar way by defining two vectors: P_k^1 defined as above and Q_k^1 representing sensor responses of the kth sensor of a second device under the illuminant E^1. By substituting these vectors in Equation (9) we can measure the degree of rank correlation. Finally we can investigate rank order invariance across device and illumination together by comparing, for example, the vectors P_k^2 and Q_k^1.

We conducted such an analysis for a variety of imaging devices and illuminants, taking as our set of surfaces, a set of 462 Munsell chips [17]. We analysed 16 different lights, including a range of daylight illuminants, Planckian blackbody radiators, and fluorescent lights. For devices we used the spectral sensitivities of the human colour matching functions [17] as well as four digital still cameras and a flatbed scanner. Table 1 summarises the results for three conditions: the first 5 rows correspond to the case in which sensor is fixed and illumination is allowed to change, the next three to the case in which illumination is fixed and device is changed, and the last row to the case in which both illumination and device change together. using the measure defined by Equation (9). Rank correlation is shown for each device averaged over all 16 illuminants (case 1), for three illuminants (daylight, fluorescent, and tungsten) averaged over all devices (case 2), and for all devices and illuminants (case 3). In all cases, the results show a very high degree of correlation: average correlation never falls below 0.964 which represents a high degree of correlation. Minimum correlation over all devices and illuminants was 0.9303 for the 1st sensor, 0.9206, for the 2nd sensor and 0.8525 for the 3rd. Thus on the basis of these results we conclude that rank orderings are preserved to a very good approximation across a change in either or both, device and illumination. It remains to determine how we are to exploit this rank invariance in practice: this we discuss in the next section.

	1st Sensor	2nd Sensor	3rd Sensor
Across Illumination			
CMFs	0.9957	0.9922	0.9992
Cam 1	0.9983	0.9984	0.9974
Cam 2	0.9978	0.9938	0.9933
Cam 3	0.9979	0.9984	0.9972
Cam 4	0.9981	0.9991	0.9994
Scanner	0.9975	0.9989	0.9995
Across Devices			
D65	0.9877	0.9934	0.9831
cwf	0.9931	0.9900	0.9710
A	0.9936	0.9814	0.9640
Across Device and Illuminant			
	0.9901	0.9886	0.9774

Table 1: Spearman's Rank Correlation Coefficient for each sensor of a range of devices. Results are averaged over all pairs of a set of 16 illuminants.

4. Exploiting Invariance: Histogram Equalisation

There are a number of ways we might employ rank ordering information to derive an invariant representation, we set forth one such method here which we will demonstrate has a number of interesting properties. To understand our method consider a single channel of an *RGB* image recorded under an illuminant o where without loss of generality we restrict the range of R^o to be on some finite interval: $R^o \in [0 \ldots R_{max}]$. Now, consider further a value

$R_i^o \in [0 \ldots R_{max}]$ where R_i^o is not necessarily the value of any pixel in the image. Let us define by $P(R^o < R_i^o)$, the number of pixels in an image with a value less than or equal to R_i^o. Under a second illuminant, c, a pixel value R^o under illuminant o is mapped to a corresponding value R^c. We denote by $P(R^c < R_i^c)$ the number of pixel values in the second image whose value is less than R_i^c. Assuming that the illumination change preserves rank ordering of pixels we have the following relation:

$$P(R^c < R_i^c) = P(R^o < R_i^o) \tag{10}$$

That is, a change in illumination preserves cumulative proportions. Given this we define one channel of the invariant image representation thus:

$$R_i^{inv} = \frac{R_{max}}{N_{pix}} P(R^o \leq R_i^o) = \frac{R_{max}}{N_{pix}} P(R^c \leq R_i^c) \tag{11}$$

where N_{pix} is the number of pixels and the constant $\frac{R_{max}}{N_{pix}}$ ensures that the invariant image has the same range of values as the input image. Repeating the procedure for each channel of a colour image results in the required invariant image.

The reader familiar with the image processing literature might recognise Equation (11). Indeed this transformation of image data is one of the simplest and most widely used methods for image enhancement and is commonly known as histogram equalisation. Histogram equalisation is an image enhancement technique originally developed for a single channel, or grey-scale, image. The aim is to increase the overall contrast in the image since doing so typical brightens dark areas of an image, increasing the detail in those regions which in turn can sometimes result in a more pleasing image. This aim is achieved by transforming an image such that the histogram of the transformed image is as close as possible to a uniform histogram. The approach is justified on the grounds that amongst all possible histograms, a uniformly distributed histogram has maximum entropy. Maximising the entropy of a distribution maximises its information and thus histogram equalising an image maximises the information content of the output image. Accepting the theory, to histogram equalise an image we must transform the image such that the resulting image histogram is uniform. Now, suppose that x_i represents a pixel value in the original image and x_i^t its corresponding value in the transformed image. Let us further assume that x_i and x_i^t are continuous variables and let us denote by $p(x)$ and $p_t(x^t)$ the probability density functions of the original and transformed image. We would like to transform the original image such that the proportion of pixels less than x_i^t in the transformed image is equal to the proportion of image pixels less than x_i in the original image, and that moreover the histogram of the output image is uniform. This implies:

$$\int_0^{x_i} p(x)dx = \int_0^{x_i^t} p_t(x^t)dx^t = \frac{N_{pix}}{x_{max}} \int_0^{x_i^t} dx^t \tag{12}$$

Evaluating the right-hand integral we obtain and rearranging terms we have:

$$x_i^t = \frac{x_{max}}{N_{pix}} \int_0^{x_i} p(x)dx \tag{13}$$

Equation (13) tells us that to histogram equalise an image we transform pixel values such that a value x_i in the original image is replaced by the proportion of pixels in the original image which are less than or equal to x_i. A comparison of Equation (11) and Equation (13) reveals that, disregarding notation, they are the same. So, the invariant image is obtained by simply histogram equalising each of the channels of our original image. In practice, applying the histogram equalisation procedure to an image results in a transformed image whose resulting histogram is only approximately uniform. This is because the range of values a pixel can take is discrete not continuous as we assumed in the analysis above.

In the context of image enhancement it is argued [16] that applying an equalisation to the channels of a colour image separately is inappropriate since this can produce significant colour shifts in the transformed image. But in our context we are interested not in the visual quality of the image but in obtaining a representation which is illuminant and/or device invariant. Histogram equalisation achieves just this provided that the rank ordering of sensor responses is itself invariant to such changes.

5. An Application to Colour Indexing

To test the invariance properties of histogram equalisation we applied the method to an image retrieval task. Finlayson *et al* [3] recently found that existing invariant approaches were unable to facilitate good enough image indexing across a change in either, or both illumination and device. Here we repeat their experiment but using histogram equalised images as our basis for indexing to investigate what improvement, if any, the method brings.

The experiment is based on a database of images of coloured textures captured under a range of illuminants and devices and described in [4]. In summary there are 28 different coloured textures each captured under six different devices (4 cameras and 2 scanners). In addition each camera was used to capture each of the textures under 3 different lights so that in total there are $(3 \times 4 + 2) \times 28 = 392$ images. In image indexing terms this is a relatively small database and it is chosen because it allows us to investigate performance across a change in illumination and device. In our experiment we tested indexing performance

Model	Case (1)	Case (2)	Case (3)
RGB	63.23	71.85	65.53
Greyworld	93.96	94.22	92.28
Hist Eq.	96.72	95.52	94.54

Table 2: Average Match Percentile results of the indexing experiment for four different cases: (1) Across illumination, (2) Across cameras, (3) Across all devices and illumination.

across three different conditions: **(1)** across illumination, **(2)** across homogeneous devices (across cameras), and **(3)** across both devices and illumination. In each case the experimental procedure was as follows. First, we choose a set of 28 images all captured under the same conditions (same device and illuminant) to be our image database. Next we select from the remaining set of images a subset of appropriate query images. So, if we were testing performance across illumination, we selected as query images the 56 images captured by the device corresponding to the database images, under the two non-database illuminants. Then, for all database and query images we derived an invariant image using the histogram equalisation method. We then represented the invariant image by its colour distribution: that is, by a histogram of the pixel values in the invariant image. All results reported here are based on 3-dimensional histograms of dimension $16 \times 16 \times 16$.

Indexing is performed for a query image by comparing its histogram to each of the histograms of the database images. The database image whose histogram is most similar to the query histogram is retrieved as a match to the query image. We compare histograms using the intersection method described by Swain *et al* [15] which we found to give the best results on average. Indexing performance is measured using average match percentile [15] which gives a value between 0 and 100%. A value of 99% implies that the correct image is ranked amongst the top 1% of images whilst a value of 50% corresponds to the performance we would achieve using random matching. Table 2 summarises the average match percentile results for the three different conditions. In addition to results for histogram equalisation we also show results based on histograms of the original images (*RGB*), and on Grey-world normalised images (the best performing invariant representation in the original experiment [3]). Significantly, histogram equalisation outperforms Grey-world for all conditions. Histogram equalisation results across a change in illumination are very good: an AMP of close to 97% as compared to 94% for Grey-world. In absolute terms results for matching across devices (Case (2)) and across device and illuminant (Case 3) are less good. However, histogram equalisation still performs significantly better than Grey-world.

These results are good in the sense that they show the new method to outperform previous invariant approaches but they also raise a number of issues worthy of further comment. First, it is surprising that one of the simplest invariants – Grey-world – performs as well as it does, whilst other more sophisticated invariants (evaluated in the original experiment [3]) perform very poorly. This performance indicates that for this dataset a diagonal scaling of sensor responses accounts for most of the change that occurs when illuminant or device is changed. It also suggests that any non-linear transform applied to the device responses post-capture must be very similar for all devices: most likely a simple power function is applied. Secondly, given the analysis in § 3.1 we might have expected that histogram equalisation would have performed somewhat better than it does. Possible reasons for this imperfect performance can be found by an examination of the images which make up the database. In addition to differences due to device and illumination, images in the database also differ spatially: i.e. the illumination varies spatially across the extent of an image and this spatial variation differs from image to image. Images of the same scene under a constant and spatially varying illuminant do not look the same after histogram equalisation. We are currently investigating how this spatial aspect of illumination can be dealt with. Additional analysis of the indexing results also revealed that poor performance was restricted to a small number of images. An inspection of these images revealed a number of artifacts of the imaging process. Specifically a number of images captured under tungsten illumination have values of zero in the blue channel for many pixels. Further, some images with uniform backgrounds were found to have significant non-uniformities in these regions when captured with the scanners. In both cases the resulting histogram equalised images are far from invariant. Excluding these images leads to a significant improvement in indexing performance. However, for an invariant image representation to be of practical use in an uncalibrated environment it must be robust to the limitations of the imaging process. Thus we have reported results including all images. Further testing on more diverse and larger images databases is required to properly determine the power of this method as compared to other invariant approaches.

6. Conclusions

We have shown in this paper that under certain theoretical conditions the rank orderings of sensor responses are invariant both to a change in illumination and to a change in imaging device. We exploited this fact using the technique of histogram equalisation to derive a novel image representation which is, in theory, both device and illuminant invariant. We have demonstrated that in practice only quasi-invariance is achieved using this method but we have shown that the degree of invariance is sufficient to be of practical use in a set of image indexing experiments using images captured under a range of different lights with a number of different imaging devices.

References

[1] J. Bach, C. Fuller, A. Gupta, A. Hampapur, B. Horowitz, R. Humphrey, and R. Jain. The virage image search engine: An open framework for image management. In *SPIE Conf. on Storage and Retrieval for Image and Video Databases*, volume 2670, pages 76–87, 1996.

[2] James L. Dannemiller. Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vision Research*, 33(1):131–140, 1993.

[3] G. Finlayson and G. Schaefer. Colour indexing arcoss devices and viewing conditions. In *2nd International Workshop on Content-Based Multimedia Indexing*, 2001.

[4] G. Finlayson, G. Schaefer, and G. Y. Tian. The UEA uncalibrated colour image database. Technical Report SYS-C00-07, School of Information Systems, University of East Anglia, Norwich, United Kingdom, 2000.

[5] G.D. Finlayson, S.S. Chatterjee, and B.V. Funt. Color angular indexing. In *The Fourth European Conference on Computer Vision (Vol II)*, pages 16–27. European Vision Society, 1996.

[6] G.D. Finlayson, B. Schiele, and J. Crowley. Comprehensive colour image normalization. In *eccv98*, pages 475–490, 1998.

[7] Graham Finlayson. *Coefficient Colour Constancy*. PhD thesis, Simon Fraser University, 1995.

[8] Brian Funt, Kobus Barnard, and Lindsay Martin. Is machine colour constancy good enough? In *5th European Conference on Computer Vision*, pages 455–459. Springer, June 1998.

[9] Brian V. Funt and Graham D. Finlayson. Color Constant Color Indexing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):522–529, 1995.

[10] T. Gevers and A.W.M. Smeulders. Color based object recognition. *Pattern Recognition*, 32:453–464, 1999.

[11] Graham. D. Finlayson, Steven Hordley, and Paul Hubel. Illuminant estimation for object recognition. *COLOR research and application*, 2002. to appear.

[12] W. Niblack and R. Barber. The QBIC project: Querying images by content using color, texture and shape. In *Storage and Retrieval for Image and Video Databases I, volume 1908 of SPIE Proceedings Series*. 1993.

[13] C. Poynton. The rehabilitation of gamma. In *SPIE Conf. Human Vision and Electronic Imaging III*, volume 3299, pages 232–249, 1998.

[14] M. Stricker and M. Orengo. Similarity of color images. In *SPIE Conf. on Storage and Retrieval for Image and Video Databases III*, volume 2420, pages 381–392, 1995.

[15] Michael J. Swain and Dana H. Ballard. Color Indexing. *International Journal of Computer Vision*, 7(1):11–32, 1991.

[16] Alan Watt and Fabio Policarpo. *The Computer Image*. Addison-Wesley, 1997.

[17] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. New York:Wiley, 2nd edition, 1982.