

Stereo retinex

Weihua Xiong, Brian Funt *

School of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby, BC, Canada V5A 1S6

Received 2 February 2007; accepted 15 November 2007

Abstract

The retinex algorithm for lightness and color constancy is extended to include 3-dimensional spatial information reconstructed from a stereo image. A key aspect of traditional retinex is that, within each color channel, it makes local spatial comparisons of intensity. In particular, intensity ratios are computed between neighboring spatial locations, retinex assumes that a large ratio indicates a change in surface reflectance, not a change in incident illumination; however, this assumption is often violated in 3-dimensional scenes, where an abrupt change in surface orientation can lead to a significant change in illumination. In this paper, retinex is modified to use the 3-dimensional edge information derived from stereo images. The edge map is used so that spatial comparisons are only made between locations lying on approximately the same plane in 3-dimensions. Experiments on real images show this method works well, however, they also reveal that it can lead to isolated regions, which, as a result of being isolated, are incorrectly determined to be grey. To overcome this problem, stereo retinex is extended to allow information that is orthogonal to the space of possible illuminants to propagate across changes in surface orientation. This is accomplished by transforming the original RGB image data into a color space based on coordinates of luminance, illumination and reflectance. This coordinate system allows stereo retinex to propagate reflectance information across changes in surface orientation, while at the same time inhibiting the propagation of potentially invalid illumination information. The stereo retinex algorithm builds upon the multi-resolution implementation of retinex known as McCann99. Experiments on synthetic and real images show that stereo retinex performs significantly better than unmodified McCann99 retinex when evaluated in terms of the accuracy with which correct surface object colors are estimated.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Although it is well established that for human subjects that a surface's perceived spatial location affects the perception of its lightness and color [1,2], many machine color constancy models [3–14] make no use of 3-dimensional spatial information. In fact, many of the methods are based on binarized color histograms, which discard all the images' spatial structure, and rely instead on statistical properties of the color distribution in order to determine the color of the scene illuminant. Although these methods work quite well [5], they all assume implicitly that there is a single scene illuminant. However, multiple illuminants are common in typical scenes. Outdoors, for example, shadowed

areas are not only darker, but much bluer, than those in the sun, because the sky's light is bluer than the sun's.

In this paper, we extend retinex to take advantage of 3-dimensional distance information extracted from stereo imagery. In particular, since an abrupt change in surface orientation may lead to an abrupt change in the incident illumination as, for example, occurs due to self-shadowing, retinex is modified so that its computation does not cross edges in the depth map. In this way, it can provide lightness/color estimates for different parts of the scene that may be illuminated differently.

Although this modification of retinex does ameliorate many of problems that arise in multi-illuminant scenes, the processing has a tendency to result in isolated grey areas. This problem arises especially for surfaces of uniform color that are completely isolated from other surfaces by a change in surface orientation. Retinex normalizes to white, so any completely isolated single color will always

* Corresponding author.

E-mail address: funt@sfu.ca (B. Funt).

be made white (or grey after subsequent intensity adjustment). To overcome this problem, a color coordinate system [27] is used with axes representing variation in illumination color, intensity, and object reflectance. Retinex is applied separately to each of these color channels and the result is then transformed back to the original color coordinates. The coordinate system allows stereo retinex to propagate reflectance information across changes in surface orientation, while at the same time inhibiting the propagation of potentially invalid illumination information.

Tests on synthetic and real images show that the modified, depth-aware stereo retinex method outperforms the original retinex method in terms of the accuracy with which the true scene surface colors are estimated. Accurate estimation of scene colors under uncontrolled illumination conditions is important in many computer vision applications.

1.1. Retinex background

Retinex has a long history beginning with an early paper by Land and McCann [15] and there are many variations on the original retinex algorithm. The basic principles of retinex are: (1) color is obtained from 3 ‘lightnesses’ computed separately for each of the color channels; (2) the ratios of intensities from neighboring locations are assumed to be illumination invariant; (3) lightness in a given channel is computed over large regions based on combining evidence from local ratios; (4) the location with the highest lightness in each channel is assumed to have 100% reflectance within that channel’s band. Lightness refers to the perceived (in the case of human perception), or estimated (in the case of computational methods) surface albedo (reflectance averaged over the channel’s band).

The initial versions of retinex were based on combining the ratio information along random paths across the image. Multi-resolution versions of retinex were introduced for efficiency [16]. Horn [17] formalized retinex in terms of differentiation, thresholding and re-integration in the logarithm domain. Kimmel et al. [18] formulate the computation as a variational optimization problem. Two versions of retinex have been given standardized definitions in terms of Matlab code [19].

All of the retinex variants treat the input image as a spatial arrangement of colors and make no use of the 3-dimensional structure of the underlying scene. However, there are a number of psychophysical experiments indicating that the human lightness and color perception are influenced by information from several sources, including 3-dimensional scene geometry. In particular, Gilchrist’s early experiments [20] showed that, in the black and white scenes, changing a surface’s apparent 3-dimensional context affected the perception of its lightness. Gilchrist writes, “The central conclusion of this research is that perceived surface lightness depends on ratios between regions perceived to lie next to one another in the same plane” [20]. The extension to retinex proposed here uses ratios between

regions lying next to one another and, furthermore, specifically excludes ratios from neighboring regions lying in different planes. In experiments using computer graphics rendered 3-dimensional scenes, Boyaci et al. [21] provided further evidence for the relationship between perceived orientation and the perceived lightness of matte surfaces. Yamauchi et al. [22] used stereoscopic stimuli to support the notion that surface color perception is strongly influenced by depth information. Bloj et al. [23] illustrated the effect of spatial shape on chromatic recognition. Yang and Shevell [24] show that binocular disparity can improve color constancy. Adelson [25] argues that statistical and spatial arrangement information are combined for lightness perception.

Since there is plenty of psychophysical evidence indicating a connection between a surface’s spatial properties in 3-dimensions and its perceived lightness and color properties, the question is how to include the spatial information into a color constancy model? We investigate how it can be incorporated into the retinex model in particular, and show that spatial information does improve its color constancy performance significantly.

2. Stereo retinex

Since we begin with the multi-resolution version of the retinex algorithm, known as McCann99 [19], and extend it to include 3D spatial information, we briefly describe the original algorithm. McCann99 is a multi-resolution technique which involves the standard pyramid of decreasing resolution. The computation starts at the top of the pyramid with a ratio-product-reset-average process that involves local comparisons between each pixel and its immediate neighbors. The procedure is iterative so that a pixel’s lightness estimate is updated based on its current lightness estimate in conjunction with its intensity ratios with respect to its neighbors. After a fixed, but user-selectable, number of iterations, the lightness estimates are propagated down a layer where the computation is continued, then propagated further.

We use a stereo image to calculate a depth map registered with the image data. Details of the camera setup, calibration and stereo-correspondence algorithm will be described in the Experiments section; however, any standard stereo-reconstruction algorithm could be used. Edges in the depth map are then detected using a modified version of the method proposed by Gelautz et al. [26]. These edges represent sharp changes in surface orientation, or depth discontinuities such as those created by occlusion.

The depth edges are the key factor in controlling the spatial comparisons made during the retinex computation. Traditional retinex compares a pixel to all its neighbors. In this case, the implicit assumption is that a large change in intensity between pixels arises from a change in surface reflectance, while a small change arises from a gradual change in illumination. However, in 3-dimensions an abrupt change in surface orientation can also mean that

there is an abrupt change in the incident illumination, since the differently oriented parts of the surface may be pointed towards different light sources. Similarly, depth discontinuities imply that there are two separate surfaces, which may, of course, be illuminated differently. As Gilchrist [20] pointed out, the only spatial comparisons between neighboring locations on the same locally planar surface should be used. With the additional information about the location of depth edges derived from stereo, the proposed stereo retinex method only makes comparison between pixels that do not cross a depth edge. Although this is conceptually simple, the computation requires some organization, especially to accommodate the multi-resolution aspect of McCann99 retinex.

Since McCann99 retinex compares values at neighboring pixels and averages lightness estimates from them as well, what is required is an efficient way to stop it making comparisons across depth edges. This is accomplished by first constructing separate maps for vertical and horizontal edges elements. This division makes it easier to propagate the edges up to lower resolution levels of the multi-resolution pyramid. Once the edge information is propagated through the pyramid, a bit-mask is used to encode the subset of the immediate a pixel's 8 neighbors that are all on the same side of any edges. As McCann99 iterates, it simply uses the bit-mask encoding to determine which neighbors to visit. Details are given below in the "Implementation Details" section.

3. Stereo retinex in LIS color coordinates

Fig. 1 demonstrates a problem that can arise with stereo retinex when spatial edges isolate regions from one another. If all spatial comparison across the edge is inhibited then the color information will not propagate at all. In this case, some areas will tend to become grey. This problem becomes especially acute for surfaces of uniform color that are completely isolated by spatial edges. Because retinex normalizes to white, any completely isolated single color will always become grey. The final result is grey, not white, because in the all the figures below, a pixel's output intensity is made to match its input intensity. The synthetic scene in Fig. 1 is composed of two patches meeting at a

sharp angle. There is tungsten illumination illuminating the blue patch from left, while D65 is illuminating the red patch from the right. For stereo retinex, the spatial edge between them isolates them from one another, so both turn grey.

To mitigate against this greying problem, we use a color coordinate system that will allow retinex to pass information about surface reflectance across 3D orientation changes while still inhibiting the exchange of possibly incorrect illumination information. The goal of the coordinate system is to represent illumination change, reflectance and luminance in as independent components as possible. Without some additional assumptions, this would not be possible. However, it can be done to a certain extent by exploiting the 1-dimensional constraint on illumination [27].

We model the RGB sensor response in the standard way,

$$p_k = \int E(\lambda)S(\lambda)R_k(\lambda) \quad k = R, G, B \quad (1)$$

$E(\lambda)$, $S(\lambda)$, $R(\lambda)$ are the illumination spectral power distribution, matte surface reflectance function, and sensor sensitivity, respectively. If the sensor sensitivities are narrow band, they can be modeled as Dirac delta functions and (1) reduces to,

$$p_k = E(\lambda_k)S(\lambda_k) \quad (2)$$

Following [27], let us further suppose that the illumination can be approximated as a blackbody radiator described by Planck's law,

$$E(\lambda, T) = I c_1 \lambda^{-5} e^{-\frac{c_2}{\lambda T}} \quad (3)$$

I is the power of illumination, T is the blackbody radiator temperature, and the constants C_1 and C_2 are $3.74183 \cdot 10^{-16} \text{ Wm}^2$ and $1.4388 \cdot 10^{-2} \text{ mK}$, respectively. Eq. (2) becomes

$$p_k = I c_1 \lambda^{-5} e^{-\frac{c_2}{\lambda T}} S(\lambda_k) \quad (4)$$

Taking logarithms, we have [27]

$$\log(p_k) = \log I + \log(S(\lambda_k)) - \frac{c_2}{T \lambda_k} + \log(c_1 \lambda_k^{-5}) \quad (5)$$




Fig. 1. (a) A synthetic scene composed of two patches. The blue one is lit by tungsten light from the left; the red one is lit by D65 from the right. (b) The image (monocular version) input to stereo retinex. The red line is the spatial edge between them, inserted manually in this case. (c) Both patches appear gray after stereo retinex because they are isolated surfaces.

In this equation, $\log I$ relates to luminance; $\log(S(\lambda_k))$ depends only on surface reflectance; $c_2/(T\lambda_k)$ depends only on illumination; and the last term is a constant determined by the sensor sensitivity. After combining the equations and canceling out the terms $\log I$ and $-\frac{C_2}{\lambda_k}$, we have

$$\begin{aligned} & \left(1 + \frac{n_R}{n_G - n_R} - \frac{n_B}{n_G - n_B}\right) \log R - \frac{n_B - n_R}{n_G - n_R} \log G - \log B \\ & = m_R + \frac{n_R m_G - n_R m_R + n_B m_B}{n_G - n_R} - m_B \end{aligned} \quad (6)$$

The n_k are fixed by the choice of camera sensitivity, and m_k are fixed by the choice of camera and surface reflectance. Therefore, for a given surface reflectance, varying the illumination’s color temperature or its luminance causes $(\log R, \log G, \log B)$ to move within a plane. The planes generated for different surface reflectances are all parallel to one another.

Fig. 2 shows the planes formed by 3 sample surfaces under the 102 illuminant spectra from the Simon Fraser University database [28] at 15 different luminance values each. These 102 illuminants are not specifically black-body radiators, but common light sources found around a university campus; nevertheless, the planar model works well. PCA (principal component analysis) determines the plane and establishes that the first 2 dimensions explain 99.1% percent of the variance. The PCA axes define the color coordinate system to be used by stereo retinex. We label the resulting axes as L, I and S for Luminance, Illumination and Surface, and the space the LIS color space.

The basic stereo retinex method described above is modified so that, at a 3D surface edge, information is allowed to propagate within the channel representing surface reflectance, while it continues to be inhibited within the illumination and intensity channels.

4. Implementation details

The main difficulty in implementing stereo retinex as a modification of the McCann99 algorithm is in transmitting the spatial edge information from one level of the multi-resolution pyramid to the next. For convenience, the edges found from the stereo depth map are assumed to lie in between image pixels. To propagate the edge information to the next lower resolution level in the pyramid, the rewrite rules shown in Fig. 3 are used. For a 2-by-2 group of pixels, if they are all to one side of an edge then the edge is easily propagated to the next level. For the case where a vertical edge runs through the group, it is randomly assigned to pass on one side of the group or the other; or above or below it in the case of a horizontal edge.

If there are any edges between a pixel and its neighbors, then it should only make comparisons with a subset of those neighbors. This subset is compactly represented by the ‘on’ bits in an 8-bit mask using 1 bit for each of a pixel’s 8 immediate neighbors. This strategy is useful for reducing the memory requirements. Deciding whether or not an edge must be crossed to reach a neighbor to the east, south, west, or north is straightforward because the edges are either above or to the side of a pixel. For a diagonal neighbor, the one to the northeast for example, an edge must be crossed if there are edges both to the north and to the east. Together they surround the pixel’s northeast corner forming an edge as shown in Fig. 4a. Similarly, an edge must be crossed to reach either of the 2 shaded pixels in Fig. 4b.

At each iteration, McCann99 compares each pixel to its neighbors and averages the local lightness estimates. The algorithm is modified to use the 8-bit neighbor mask to indicate what subset of the neighbors to use. The number of ‘on’ bits also indicates the number to divide by in the averaging step.




Fig. 2. $(\log R, \log G, \log B)$ obtained from three different surface reflectances under 102 illuminants at 15 various intensities. Each surface is plotted with a different color. Each set lies close to a plane, and the planes corresponding to the different surfaces are parallel. The three colored lines indicate the LIS coordinate system.




Fig. 3. Rewrite rules using in propagating edge information to the next lower resolution. An edge running through the middle of a 2-by-2 region is randomly assigned to one side or the other. Vertical edges are shown here. Horizontal edges are treated analogously.




Fig. 4. (a) From the center pixel, the three shaded pixels in the upper right cannot be reached without crossing an edge. (b) The two pixels that cannot be reached are shaded.

For stereo matching, we use a fast cross-correlation, rectangular sub-regioning, and 3D maximum-surface techniques in a coarse-to-fine scheme [29]. However, noise in the image, as well as errors in calibration and rectification, can lead to false matches being made that lead to errors in the depth map. To improve the accuracy of detected spatial edges, we use the ‘edge combination’ technique developed by Gelautz et al. [26]. We used their original method with the exception of using Laplacian of Gaussian edge detection in place of Canny detection, since for our purposes it seemed to give slightly better results.

5. Experiments

We implemented stereo retinex in Matlab 7.0 by downloading and modifying the McCann99 Matlab code available from the Simon Fraser University Computational Vision Laboratory [28]. We then tested it on both synthetic and real images. Retinex’s performance is evaluated in terms of the accuracy with which it estimates the chromaticity of surface colors as they would occur under a canonical ‘white’ illumination.

Images were captured using a Kodak DCS460 single-lens reflex digital camera. A “LOREO 3D lens in a cap” is attached in place of the standard lens so that the camera records a stereo pair within a single image frame [30]. Camera geometry calibration, image rectification and stereo matching were conducted using standard procedures [29,31]. We use the stereo image to calculate a 3D depth map and then detect edges in the depth map using a modified version of the method proposed by Gelautz et al. [26].

We evaluate performance in terms of the distance between colors in rg -chromaticity ($r = R/(R + G + B)$, $g = G/(R + G + B)$) space, and in terms of the angle between colors viewed as vectors in RGB space. These are given by the following formulas, where subscript ‘ e ’ indicates the result of retinex processing, and ‘ w ’ indicates the ‘benchmark’ color under white light:

$$Ed = \sqrt{(r_e - r_w)^2 + (g_e - g_w)^2} \quad (7)$$

$$Ea = \cos^{-1} \left[\frac{(r_e, g_e, b_e) \circ (r_w, g_w, b_w)}{\sqrt{r_e^2 + g_e^2 + b_e^2} * \sqrt{r_w^2 + g_w^2 + b_w^2}} \right] \quad (8)$$

We report four basic statistical measures of the error distributions: mean, median, RMS (root mean square) and M_{\max} . M_{\max} is the average value of the largest p percent of the errors. M_{\max} is more stable with respect to presence of an isolated extreme value than the simple maximum. In this paper, p is set to be 0.5. Hordley et al. [32] indicate that the median angular error is often the most appropriate one to use when evaluating color constancy. RMS of the errors from N pixels is given by the standard formula:

$$RMS = \frac{1}{N} \sqrt{\sum_{i=1}^N E_i} \quad (9)$$

5.1. Tests using synthetic images

Since stereo reconstruction and edge detection will be imperfect, one goal of the synthetic-image tests is to determine how much undetected edges will affect accuracy. It is also useful to compare the performance of stereo retinex to McCann99 retinex in a controlled, noise-free environment, with completely accurate ground-truth data.

The synthetic images are constructed with a variable number of patches of different reflectances selected from the 1995 reflectances available in the database described by Barnard [28]. The illumination spectrum and sensor sensitivity functions [28] of a SONY DXC-930 3-CCD camera are used to derive the RGB for each patch. First, a benchmark image is generated using ‘white’ illumination. Second, using the same patch reflectances, the same synthetic scene is divided into two parts. RGB’s for one part are synthesized using the spectrum of tungsten light, and for the other using D65 daylight. All the reflectance and illuminant data were downloaded from the Simon Fraser University color database [28]. For the synthetic case, we do not synthesize stereo images, but instead create the depth-edge map manually so that the number and extent of leaks between the two differently illuminated parts of the image can be controlled.

For the first experiment, we divided the image down the middle. We apply stereo retinex to the image once provid-




Fig. 5. Results for synthetic images containing only a single edge down the middle of the image. The illumination on the left half is tungsten, and on the right D65. The black line indicates the manually defined edge (a) Input image; (b) the benchmark image; (c) standard McCann99 applied in log RGB space; (d) stereo retinex applied using log RGB space; (e) McCann99 result applied using the LIS color channels; (f) stereo retinex applied using the LIS color channels with 3D edge information inhibiting propagation only within the illumination and intensity channels.

ing it a perfect edge map, and then a second time with an edge map containing a single leak. The results are shown visually in Fig. 5 and tabulated numerically in Table 1.

For the second experiment, the image is separated into 2 parts via an irregular border. The irregular border tests the effectiveness of the propagation of the edge information through the multi-resolution pyramid. The results are shown in Fig. 6 and Table 1.

5.2. Tests using real images

We conducted two sets of experiments with real images. In the first, the only objects in the scene were Macbeth ColorCheckers [33]. In the second, other more typical objects were included. Although scenes such as a room with tungsten light from a lamp along with daylight from a window are common, we arranged a controlled 2-illuminant environment. Two tungsten lamps were used with filters




Fig. 6. Irregular boundary between the two regions. The edge separating the regions is defined manually. (a) Input image; (b) the benchmark image; (c) standard McCann99 applied in log RGB space; (d) stereo retinex applied using log RGB Space; (e) McCann99 result applied using the LIS channels; (f) stereo retinex applied using the LIS channels.

Table 1

Performance comparison of the synthetic image cases from Fig. 5 with straight edge boundary, and Fig. 6 with an irregular edge boundary of SR LIS (stereo retinex processed using LIS color channels); SR (stereo retinex processed using log RGB space), M99 LIS McCann 99 retinex processed in LIS color channels; and M99 (McCann99 retinex processed using log RGB space)

		Distance ($\times 10^2$)				Angular			
		M_{\max}	RMS	Mean	Median	M_{\max}	RMS	Mean	Median
Straight boundary	SR LIS	8.16	3.44	2.97	2.54	4.67	3.51	3.13	2.86
	SR	10.43	4.05	3.54	3.21	5.01	4.49	4.01	3.78
	M99 LIS	11.61	5.01	4.37	3.83	11.94	5.32	4.85	4.58
	M 99	10.13	4.99	4.55	4.43	11.94	5.59	5.19	5.13
Irregular boundary	SR LIS	6.52	3.08	2.72	2.62	7.06	3.06	2.73	2.61
	SR	6.69	3.53	3.26	3.21	10.19	4.08	3.70	3.47
	M99 LIS	12.27	4.83	4.22	3.58	13.27	5.11	4.62	4.33
	M99	12.63	4.97	4.33	4.12	13.88	5.19	4.98	4.65

attached. One, with a blue filter, lit the scene from the left; the other, with a red filter, lit the scene from the right.

The first scene consisted of two Macbeth ColorCheckers meeting at an angle as shown Fig. 7. The scene was then photographed in stereo. To obtain a benchmark image, a white reflectance standard was introduced at the side of the scene and then an additional image was taken using white light. The RGB channels were then scaled in order to make the reflectance standard perfectly white (i.e.,




Fig. 7. Comparison of standard retinex to stereo retinex both in log *RGB* and in LIS coordinates operating on the image of a simple scene lit with bluish light from the left and reddish light from the right. (a) Input image of a two-illuminant scene; (b) the white-point adjusted benchmark image; (c) standard McCann99 applied in log *RGB* space; (d) stereo retinex applied using log *RGB* space; (e) McCann99 result applied to LIS color channels; (f) stereo retinex applied in LIS color channels with 3D edge information inhibiting propagation only within the illumination and intensity channels.

$R = G = B = 255$). Results are shown in Fig. 8 and Table 2.

The surface orientation edge in the previous scene is very distinct and easily identified. To test how well stereo retinex works in a less controlled environment, we use the more complex scenes shown in Figs. 9 and 10. Again, Fig. 9 has blue light from the right and red light from the left. As can be seen from the white bust in the upper right, as well as the white button in the lower left, stereo retinex in log *RGB* (Fig. 9e) is more successful at eliminating the illumination variation than McCann99 (Fig. 9d). Both methods push the colors towards grey because retinex normalizes colors relative to the whitest surface within a local region. This leads to desaturation of the colors when there is no nearby white surface. In the case of stereo retinex, this problem is exacerbated by the fact that depth edges (correctly) limit the distance within which a white surface needs to be found. Using the LIS color space, more surface color information propagates across the edges and this leads to the more colorful result (Fig. 9g, Table 3).

Both the ColorChecker and toy scenes have two distinct illuminants, but even in a single-illuminant scene the illumination can vary locally due to light interreflecting off colored surfaces. Fig. 10 shows an example of a single-illuminant scene. One example of the advantage of stereo retinex over McCann99 can be seen by comparing the left facing part of the horizontal book, which is in shadow so that it is only being illuminated indirectly. In the McCann99 result, on the book cover there is a region with a pink cast as well as one with a pale green cast; whereas, stereo retinex in LIS space correctly removes the original red cast. Overall performance results are tabulated in Table 4.

6. Retinex's iteration parameter

One of the key parameter choices to make when running McCann99 retinex is the number of iterations to be conducted at each pyramid level. The larger the number of iterations the greater distance at which pixels affect one another. Fig. 11 plots the median chromaticity error as a function of the number of iterations for the scene from Fig. 10. The plots for all the other scenes showed a similar trend. From this plot, it appears that 1 iteration is the best




Fig. 8. Edge map and recovered illumination: (a) edges representing abrupt changes in surface orientation extracted from the stereo image pair are marked in white; (b) chromaticity of illumination as estimated by stereo retinex in LIS color channels correctly shows a sharp change in illumination where the surface orientation changes; (c) Illumination field recovered by McCann99 shows a much less distinct change in illumination.

Table 2

Two-illuminant real image performance comparison of SR LIS (stereo retinex processed using LIS color channels), SR (stereo retinex processed in log RGB space), M99 LIS McCann 99 retinex processed in LIS color channels, and M99 (McCann99 retinex processed in log RGB space)

	Distance ($\times 10^2$)				Angular			
	M_{\max}	RMS	Mean	Median	M_{\max}	RMS	Mean	Median
SR LIS	8.51	2.98	2.77	2.68	8.88	3.80	3.63	2.99
SR	9.44	3.09	2.79	2.73	9.72	4.06	3.67	3.59
M99 LIS	16.93	5.32	3.78	3.83	13.62	7.14	5.46	4.38
M99	18.99	5.81	4.77	3.89	16.73	7.69	6.25	4.77




Fig. 9. Real image performance comparison. (a) Input image of two-illuminant scene of toys with uniform background illuminated with reddish light from the left and bluish from right; (b) white-point adjusted benchmark image; (c) edge map in which the arrow indicates where edges completely isolate the toy's green tongue from all other regions; (d) standard McCann99 applied in log RGB space; (e) stereo retinex applied using log RGB Space, the isolated small patch turns gray; (f) McCann99 result applied to channels of the LIS color coordinate system; (g) stereo retinex applied in the LIS color channels with 3D edge information inhibiting propagation only within the illumination and intensity channels, the isolated small patch is close to the green it should be as in the (b). (h–k) Error maps corresponding to the results from (d–g) in which large errors are shown as dark and zero error as white.

choice, so it is what has been used to obtain all the results reported above.

7. Conclusion

The McCann99 retinex method was modified to include information about the 3-dimensional structure of the imaged scene. The additional 3-dimensional information is obtained from stereo imagery. Fundamental to retinex is that it ratios intensities from neighboring image locations. Stereo retinex specifically stops retinex from using ratios that occur across abrupt changes in 3-dimensional surface orientation, or across abrupt changes in depth. It thereby avoids abrupt changes in the incident illumination from having a deleterious effect upon its calculations. This strategy is in line with Gilchrist's experiments [20] that showed how spatial context affects human lightness perception and his conclusion that the important ratios are the

ones relating to locations lying on the same 3-space plane. Although stereo imagery was used here to determine the 3-dimensional structure, any other method (e.g. from shading in a monocular image) of identifying when neighboring image pixels correspond to scene points lying on a locally planar surface would work just as well.

Although a significant improvement over traditional retinex, stereo retinex also highlights the problem that limiting the propagation of lightness information across the image increases the likelihood that it will normalize colors relative to a color which is not a true white, with the result that some colors are estimated as being more desaturated than they should be. To solve this problem, the LIS color coordinate system was applied during retinex processing. The LIS system defines channels that relate to changes in illumination, intensity and reflectance. Both retinex and stereo retinex applied to these channels performs modestly better than when either is




Fig. 10. Real-image performance comparison. (a) Input image of single-illuminant scene of books illuminated solely by reddish light from the right; (b) the white-point adjusted benchmark image; (c) standard McCann99 applied in log RGB space; (d) stereo retinex applied using log RGB Space; (e) McCann99 applied in LIS color channels; (f) stereo retinex applied in LIS space with 3D edge information inhibiting propagation only within the illumination and intensity channels. Note how the color of the orange and yellow patches on the ball are recovered better in this case. Also the pink illumination cast is removed more completely. (g–j) Error maps corresponding to the results from (c–f) in which large errors are shown as dark, and zero error as white.

Table 3
Two-illuminant image of toys against a gray background. Performance comparison between SR LIS (stereo retinex processed using LIS color channels); SR (stereo retinex processed in log RGB space), M99 LIS McCann 99 retinex processed in LIS color channels); and M99 (McCann99 retinex processed in log RGB space)

	Distance ($\times 10^2$)				Angular			
	M_{\max}	RMS	Mean	Median	M_{\max}	RMS	Mean	Median
SR LIS	34.12	4.11	2.71	1.73	31.92	4.31	3.31	2.36
SR	43.60	7.82	4.93	3.04	39.71	7.91	5.51	3.75
M99 LIS	53.68	5.83	4.51	3.16	41.62	6.37	5.27	4.10
M99	57.18	7.02	5.46	4.13	47.25	7.73	6.40	5.32

Table 4
Single-illuminant real image books scene performance comparison between SR LIS (stereo retinex processed using LIS color channels); SR (stereo retinex processed in log RGB space), M99 LIS McCann 99 retinex processed in LIS space); and M99 (McCann99 retinex processed in log RGB space)

	Distance ($\times 10^2$)				Angular			
	M_{\max}	RMS	Mean	Median	M_{\max}	RMS	Mean	Median
SR LIS	10.64	4.03	2.92	1.82	10.72	4.24	3.12	2.02
SR	18.56	5.59	3.83	2.47	14.92	6.37	4.53	2.89
M99 LIS	24.31	6.47	4.84	3.44	24.42	6.92	5.29	3.78
M99	30.53	6.71	5.06	3.81	30.61	7.28	5.71	4.48

applied to the standard log RGB channels. By at least partially separating changes in surface reflectance from changes in illumination and intensity, the LIS color space makes it possible to express the fact that across an abrupt change in 3D surface orientation, the comparison of surface reflectance information across the edge

remains valid even though the illumination may have changed in unpredictable ways.

Stereo retinex consistently outperforms McCann99 retinex in its ability to estimate the chromaticity of surface colors as they would appear under ideal white light. For the case of retinex at least, this demonstrates that knowl-




Fig. 11. Median angular error as a function of the number of retinex's iterations parameter. The number of iterations affects the distance with which lightness information propagates across the image. Results here are for processing Fig. 10, but the trend is the same for the other scenes as well.

edge of scene's 3-dimensional spatial structure can be useful for color constancy.

Acknowledgement

This research was funded by the Natural Sciences and Engineering Research Council of Canada.

References

- [1] A.L. Gilchrist, Perceived lightness depends on perceived spatial arrangement, *Science* 195 (1977) 185–187.
- [2] J.N. Yang, L.T. Maloney, Illuminant cues in surface color perception: tests of three candidate cues, *Vision Research* 41 (2001) 2581–2600.
- [3] D. Forsyth, A novel algorithm for color constancy, *International Journal of Computer Vision* 5 (1990) 5–36.
- [4] G.D. Finlayson, Retinex viewed as a gamut mapping theory of color constancy, *Proceedings of AIC International Color Association Color 97*, vol. 2, Kyoto, Japan, 1997, pp. 527–530.
- [5] K. Barnard, L. Martin, A. Coath, B. Funt, A comparison of computational colour constancy algorithms. part two: experiments on image data, *IEEE Transactions on Image Processing* 11 (2002) 985–996.
- [6] V. Cardei, B. Funt, K. Barnard, Estimating the scene illumination chromaticity using a neural network, *Journal of the Optical Society of America A* 19 (12) (2002) 2374–2386.
- [7] B. Funt, W.H. Xiong, Estimating illumination chromaticity via support vector regression, *Proceedings of 12th Color Imaging Conference – Color Science, Systems & Applications*, 2004, pp. 47–52.
- [8] G.D. Finlayson, S.D. Hordley, P.M. Hubel, Color by correlation: a simple, unifying framework for color constancy, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23 (11) (2001) 1209–1221.
- [9] K. Barnard, L. Martin, B. Funt, Colour by correlation in a three-dimensional colour space, *Proceedings of 6th ECCV*, Dublin, 2000, pp. 275–289.
- [10] C. Rosenberg, M. Hebert, S. Thrun, Color constancy using KL-divergence, *Proceedings of 8th ICCV* 1 (2001) 239–246.
- [11] G. Sapiro, Color and illumination voting, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12 (11) (1999) 1210–1215.
- [12] E.H. Land, J.J. McCann, Lightness and retinex theory, *Journal of the Optical Society of America A* 61 (1971) 1–11.
- [13] E.H. Land, Recent advances in retinex theory, *Vision Research* 26 (1986) 7–22.
- [14] B.K.P. Horn, Determining lightness from an image, *Computer Graphics and Image Processing* 3 (1974) 277–299.
- [15] E. Land, J. McCann, Lightness and retinex theory, *Journal of the Optical Society of America A* 61 (1) (1971), January.
- [16] J. Frankle, J. McCann, Method and Apparatus for Lightness Imaging, US Patent #4, 384, 336, May 17, 1983.
- [17] B.K.P. Horn, Determining lightness from an Image, *Computer Graphics and Image Processing* 3 (1974) 277–299.
- [18] R. Kimmel, M. Elad, A. Shaked, R. Keshet, I. Sobel, A variational framework for retinex, *International Journal of Computer Vision* 52 (1) (2003) 7–23.
- [19] B. Funt, F. Ciurea, J. McCann, Retinex in matlab, *Journal of the Electronic Imaging* (2004) 48–57.
- [20] A.L. Gilchrist, Perceived lightness depends on perceived spatial arrangement, *Science* 195 (1977) 185–187.
- [21] H. Boyaci, L.T. Maloney, S. Hersh, The effect of perceived surface orientation on perceived surface albedo in binocularly viewed scenes, *Journal of Vision* 3 (2003) 541–553.
- [22] Y. Yamauchi, K. Uchikawa, Depth information affects judgment of the surface-color mode appearance, *Journal of Vision* 5 (2005) 515–523.
- [23] M.G. Bloj, D. Kersten, A.C. Hurlbert, Perception of three-dimensional shape influences colour perception through mutual illumination, *Nature* 42 (1999) 23–30.
- [24] J.N. Yang, S.K. Shevell, Stereo disparity improves color constancy, *Vision Research* 42 (2002) 1979–1989.
- [25] E.H. Adelson, *Lightness perception and lightness illusions*, second ed., New Cognitive Neuroscience, MIT Press, 2000.

- [26] M. Gelautz, D. Markovic, Recognition of object contours from stereo images: an edge combination approach, Proceedings of 2nd International Symposium on 3D Data Processing, Visualization and Transmission, pp. 774–780.
- [27] G.D. Finlayson, S.D. Hordley, Color constancy at a pixel, Journal of the Optical Society of America A 18 (2) (2001) 253–264.
- [28] K. Barnard, L. Martin, B.V. Funt, A. Coath, A data set for color research, Color Research and Application 27(3) (2002) 140–147, (Available from: <www.cs.sfu.ca/~colour>).
- [29] C.H. Sun, Fast stereo matching using rectangular subregioning and 3D maximum-surface techniques, International Journal of Computer Vision 47 (2002) 99–117.
- [30] Available from: <http://www.vision.caltech.edu/bouguetj/calib_doc/>.
- [31] Available from: <<http://www.loreo.com/>>.
- [32] S.D. Hordley, G.D. Finlayson, Reevaluation of color constancy algorithm performance, Journal of the Optical Society of America A 23 (5) (2006) 1008–1020.
- [33] Available from: <www.gretamacbeth.com>.