

Improved Colour to Greyscale via Integrability Correction

Mark S. Drew^a, David Connah^b, Graham D. Finlayson^b and Marina Bloj^c

^aSchool of Computing Science, Simon Fraser University, BC, Canada;

^bDepartment of Computer Science, University of East Anglia, UK;

^cBradford School of Optometry and Vision Sciences, School of Life Sciences, University of Bradford, UK

ABSTRACT

The classical approach to converting colour to greyscale is to code the luminance signal as a grey value image. However, the problem with this approach is that the detail at equiluminant edges vanishes, and in the worst case the greyscale reproduction of an equiluminant image is a single uniform grey value. The solution to this problem, adopted by all algorithms in the field, is to try to code colour difference (or contrast) in the greyscale image. In this paper we reconsider the Socolinsky and Wolff algorithm for colour to greyscale conversion. This algorithm, which is the most mathematically elegant, often scores well in preference experiments but can introduce artefacts which spoil the appearance of the final image. These artefacts are intrinsic to the method and stem from the underlying approach which computes a greyscale image by a) calculating approximate luminance-type derivatives for the colour image and b) re-integrating these to obtain a greyscale image. Unfortunately, the sign of the derivative vector is sometimes unknown on an equiluminant edge and, in the current theory, is set arbitrarily. However, choosing the wrong sign can lead to unnatural contrast gradients (not apparent in the colour original). Our contribution is to show how this sign problem can be ameliorated using a generalised definition of luminance and a Markov relaxation.

1. INTRODUCTION

The standard approach to converting a colour image into greyscale is to obtain the luminance value at each pixel and to use this as a grey value. This approach suffers from the drawback that regions of high colour contrast, but low luminance contrast, lose their saliency in the resulting greyscale image.¹ The limit of this problem occurs at equiluminant colour boundaries, where regions of distinct equiluminant colours may be assigned identical grey values, thus losing important image features. Indeed, such equiluminant boundaries have been found to be common in both natural and synthetic images.^{2,3}

This problem has received a large amount of recent attention,^{1,4-7} both because of the still-common usage of greyscale printing devices and an intrinsic interest in this special case of dimensionality reduction. The underlying goal of the methods that have been proposed is to create a more perceptually relevant greyscale image by preserving colour contrast as greyscale contrast in the final image. In addition to preserving image content, this approach has also been shown to generate images that are consistently preferred to luminance renderings.⁸

One of the best performing, and mathematically most elegant, solutions is the algorithm of Socolinsky and Wolff.^{9,10} The goal of this method is to code *local* colour contrast as greyscale contrast in the final greyscale image. The method works by defining the local contrast as the *gradient* of the colour image, estimating the gradient for each pixel, and then re-integrating to form a greyscale image.

Estimating the gradient of a colour image, however, is not straightforward: the derivative can be approximated in each R, G, B channel independently using local differencing, but how should the gradients be combined across the three colour channels into a single gradient estimate $[\nabla_x \nabla_y]^T$? This problem was tackled by Di Zenzo¹¹ who proposed a 2×2 “colour structure-tensor” matrix to represent the gradient. The first eigenvector \mathbf{e} of the structure tensor gives the principal direction of change in the colour image, including all colour-change information, while

Further author information: (Send correspondence to Mark S. Drew)
Mark S. Drew.: E-mail: mark@cs.sfu.ca, Telephone: +1-778-782-4277

the main eigenvalue λ_1 gives the magnitude of that change. Combined together the magnitude and direction vector give a single gradient vector $\mathbf{v} = [\nabla_x \nabla_y]^T = \sqrt{\lambda_1} \mathbf{e}$ at each pixel. The sign of this vector, however, is undefined (since the vector sense of an eigenvector is undefined): i.e., there is no way to resolve the ambiguity between $-\mathbf{v}$ and \mathbf{v} given the structure tensor alone. But this ambiguity must be addressed in order to produce a pleasing greyscale image, since choosing the sign unwisely can lead to unnatural contrast gradients that are not apparent in the colour original.

Fortunately, there is a clue as to when unnatural gradients may be present: the *curl* of the gradients ($\frac{\partial^2 f}{\partial x \partial y} - \frac{\partial^2 f}{\partial y \partial x}$) calculated at a given pixel is non-zero. Furthermore, a curl value of zero implies that the gradients are natural, and the resulting image will be artefact free. We therefore ameliorate the sign ambiguities for a gradient field in such a way as to minimise the curl value over all gradients. Here we employ a two stage approach: in the first stage we use a generalised definition of luminance, based on the Minkowski norm, to find a global transform that minimises integrability artefacts. Although this global transform will reduce sign ambiguities, it may not always resolve the ambiguities over the whole image. Therefore, in the second stage of our approach we use a local method to determine the sign based upon a Markov relaxation. By identifying gradients where there is a sign ambiguity we randomly select problem pixels and choose their sign to minimise the curl value locally, and these local changes propagate throughout the image to reduce curl globally. The process is Markovian in that error is based only on the local neighbours of a pixel, and pixel selection is random. As the process iterates, error is diffused by passing correction information to neighbouring pixels. We demonstrate several examples where using this two-stage approach results in visually improved greyscale images.

2. SOCOLINSKY AND WOLFF ALGORITHM

The starting point of this work is the contrast-preserving colour-to-greyscale algorithm of Socolinsky and Wolff.⁹ This algorithm can be divided into two processes: computing the gradient at each pixel to form contrast vector field \mathbf{v} , and then re-integrating this vector field to create a new greyscale image.

By defining a colour image as a vector-valued function which maps the 2D plane \mathbb{R}^2 onto a three dimensional colour-space \mathbb{R}^3 , the gradient of the image at a point (x_0, y_0) is given by the gradient matrix:

$$\mathbf{J} = \begin{bmatrix} \frac{\partial f_1}{\partial x} & \frac{\partial f_1}{\partial y} \\ \frac{\partial f_2}{\partial x} & \frac{\partial f_2}{\partial y} \\ \frac{\partial f_3}{\partial x} & \frac{\partial f_3}{\partial y} \end{bmatrix}, \quad (1)$$

where the function $f_i(x, y)$, which maps \mathbb{R}^2 onto \mathbb{R} , is the i th colour channel. The gradient matrix \mathbf{J} contains the partial derivatives of f in the x and y directions, and also specifies the gradient in an arbitrary direction $\mathbf{d} = [\cos \theta, \sin \theta]^T$ as $\mathbf{J} \mathbf{d}$. Assuming a Euclidean metric, the magnitude of the gradient m in an arbitrary direction is thus given by:

$$m = \mathbf{d}^T \mathbf{J}^T \mathbf{J} \mathbf{d}. \quad (2)$$

The 2×2 matrix $\mathbf{J}^T \mathbf{J}$ is Di Zenzo's structure tensor,¹¹ and determines the gradient magnitude in any direction \mathbf{d} . To obtain a single gradient estimate $\mathbf{v} = [\Delta x \ \Delta y]^T$, the goal is to find a direction \mathbf{d} for which the magnitude of the gradient m is maximal. This optimisation turns out to be an eigenvalue problem, and the solution is given by the first eigenvector of the structure tensor (i.e. the eigenvector associated with the largest eigenvalue), while the largest eigenvalue gives the square of the gradient-magnitude in the direction of maximal change.

To illustrate this, in Figure 1 we show a simple colour image and analyse the magnitude of the gradient at a step edge as a function of θ . The partial derivatives in the colour-gradient matrix \mathbf{J} are computed as forward-difference approximations $\frac{\partial f_i}{\partial x} = f_i[x+1, y] - f_i[x, y]$ and $\frac{\partial f_i}{\partial y} = f_i[x, y+1] - f_i[x, y]$; the structure tensor

is then constructed using Equation 2. From the figure we can see that the magnitude m is a sinusoidal function, which has two clear peaks at π radians and at 0 (or 2π) radians. These correspond to the vector $\pm[1, 0]^T$, which is clearly the correct direction for the gradient. The fact that there are two peaks in m , however, means that although the direction of maximal change may be clearly defined, the sign of the vector is not. This causes a serious problem, since when the vector field is re-integrated to form a greyscale image, it is the sign that determines which side of the edge will be brighter and which darker.

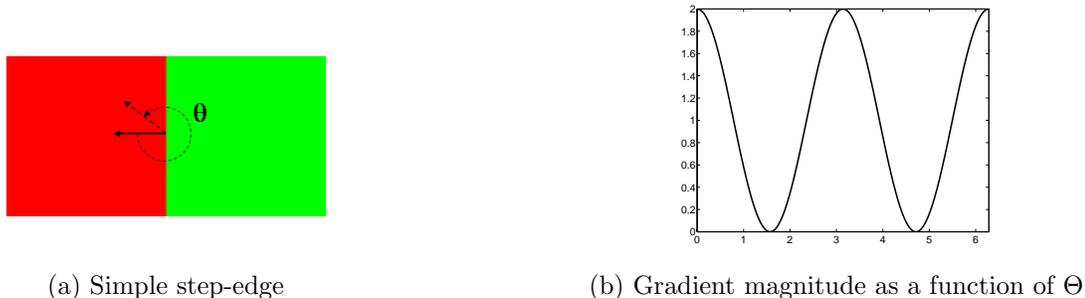


Figure 1. Gradient magnitude as a function of Θ for a simple step-edge.

One sign-assignment strategy is to appeal to the visual sensitivity of the human eye, and to take the sign of the gradient as that of the underlying luminance image, and we could use either an L_1 or L_2 -norm luminance. This approach has the desirable property that dark colours will tend to remain dark, and light colours will tend to remain light. It still has the problem, however, that the sign of the gradient is not defined for equiluminant borders. Given that the contrast along equiluminant borders is non-negligible, choosing a consistent sign assignment is important.

Fortunately it is possible to predict where such inconsistencies might occur. Assuming that the contrast-magnitudes within the image are resolvable into a valid greyscale, a consistent sign assignment will lead to the condition that:

$$\frac{\partial^2 \mathbf{f}}{\partial x \partial y} - \frac{\partial^2 \mathbf{f}}{\partial y \partial x} = 0. \quad (3)$$

where \mathbf{f} represents the original colour image, and the quantity on the left-hand-side is the *curl* of the vector field. Thus when the curl of the gradient vector field is non-zero, the field is non-integrable.¹² When regions of non-integrability coincide with equiluminant borders, this suggests the presence of inconsistent sign assignments.

3. MARKOV RELAXATION

3.1. Minkowski Luminance

In the first instance, we set the sign using a Minkowski-norm based luminance measure. We define a Minkowski-luminance image as the image whose value at each pixel is given by the Minkowski-norm for the colour vector at that pixel:

$$L_p = \left(\frac{1}{K} \sum_{i=1 \dots K} |f_i|^p \right)^{\frac{1}{p}} \quad (4)$$

where L_p is the Minkowski-luminance image, and the f_i are the individual colour channels ($K = 3$). As p increases, the p -norm approaches the maximum over the channels.¹³

The vector-sign for the colour image is then decided by pointing the eigenvector \mathbf{e} of the Di Zenzo matrix in the direction of increasing luminance. In,⁹ this approach was indeed taken, and is typically based on either an

L_1 or a standard L_2 -norm measure. However, it is quite common for images to have basically uniform luminance in either of these standard measures, particularly for graphics-generated images. For example, in the image of Monet’s painting “Impression: sunrise”, it has long been noted that the sun tends to disappear when an L_1 -type luminance is displayed (Figure 2(a)). In contrast, making use of an L_5 -based luminance brings the sun to the fore (see Figure 2(b)). We have found that the use of a Minkowski norm tends to successfully identify colour edges that are commonly hidden when using the standard norms.



(a) L_1 rendering.



(b) L_5 rendering.

Figure 2. Two L_p renderings of

Intuitively, it might make sense to choose the signs based upon an image L_p which has fewest equiluminant boundaries. In practice we repeat the sign-assignment using a range of p values, and then adopt the L_p image that leads to the lowest absolute curl value (summed over the whole vector field). Unfortunately, it can still happen that regardless of the p value that we choose, the L_p image is completely uniform, e.g. when the user has created a graphic using simple all-red, all-green etc. patches. Therefore, in such cases we replace the gradient of the L_p -luminance with a set of gradients based on the maximum over all colour channels.¹⁴

3.2. Markov Non-Integrability Reduction

We can further improve the image resulting from the approximate set of gradients derived so far by identifying problem pixels that generate integrability errors in the greyscale output. Using a Markov relaxation approach, we randomly select problem pixels and determine what local effect on the overall image integrability would result from changing the sign of the main Di Zenzo eigenvector at a pixel (thus also affecting its neighbours, and therefore propagating error-correction diffusion).

Let the putative gradient of the greyscale output be $\mathbf{v} = \{v_x, v_y\}$; each gradient component is an image-sized matrix. Integrability, as dictated by the absolute value of the curl, is not necessarily zero near pixel $\{i, j\}$, so we examine the integrability separately for each choice of sign, $\{+1, -1\}$.

The integrability (curl) around that pixel is given by $\mathcal{I} = \partial v_x / \partial y - \partial v_y / \partial x$. But, using forward-differencing as the discrete approximation of derivatives, changing the sign at $\{i, j\}$ means that \mathcal{I} is affected in the 4-neighbourhood of $\{i, j\}$. So we can form an absolute measure of integrability by examining

$$\begin{aligned}
 |\mathcal{I}| &= |(v_x[i + 1, j] - v_x[i, j]) - (v_y[i, j + 1] - v_y[i, j])| \\
 &+ |(v_x[i, j] - v_x[i - 1, j]) - (v_y[i - 1, j + 1] - v_y[i - 1, j])| \\
 &+ |(v_x[i + 1, j - 1] - v_x[i, j - 1]) - (v_y[i, j] - v_y[i, j - 1])|
 \end{aligned}$$

Thus the approach taken here is to first form a mask image M of problem pixels and then repeatedly visit them randomly (with replacement) to decide on the sign-assignment that best reduces overall integrability. To

do so, we iterate several times over complete sets of random selections of pixels. Suppose the cardinality of M is N . Then for each of a few iterations we randomly visit the set of N problem locations, in M , for a total of N times. At the end of each iteration, if the set of vector-sense signs has not altered, we exit. Finally, in case any pixels in M happen to have not been visited yet, we visit all remaining mask pixels.

In each pass, changing the integrability at pixels in M may in fact decrease the integrability for neighbouring pixels. Therefore we include *error-diffusion* by re-setting the error mask M in each main iteration, by recalculating a non-integrability image. Typically, a maximum of about 10 iterations is necessary with more usual numbers of iterations being about 3 or 4.

To set mask M , we note that problems of assignment of the vector-sense are most common when either the gradient of the luminance is small (e.g., flat regions in the image), or the product of the greyscale-gradient with the luminance gradient is small, leading to an inconclusive sign-assignment. We calculate the dot product of the main eigenvector, \mathbf{e} , with the gradient ∇L of the luminance image, and set the vector-sign according to the following scheme: Suppose the greyscale gradient has so far been determined from \mathbf{e} as \mathbf{v} , in the direction of \mathbf{e} and with magnitude given by the maximum eigenvalue of the Di Zenzo matrix. Then we determine if \mathbf{v} is pointing in the same direction as ∇L by calculating

$$\begin{aligned} S &= \mathbf{v} \cdot \nabla L; \\ \mathbf{v} &\rightarrow -\mathbf{v} \quad \text{if } S < 0 \end{aligned} \tag{5}$$

When S is small (for flat regions or equiluminant borders), the vector-sign is undetermined, so using a threshold τ_1 we set the mask of problem pixels to

$$M \equiv M_1 = (S < \tau_1) \tag{6}$$

But this indicator may over-estimate the number of problem areas, since the sign is unimportant when $\mathbf{v} [i, j] = [0 \ 0]^T$. Therefore we append another measure: problem pixels are those for which S is small and integrability error \mathcal{I} is too high:

$$M_2 = \mathcal{I} > \tau_2, \quad M = M_1 \text{ AND } M_2 \tag{7}$$

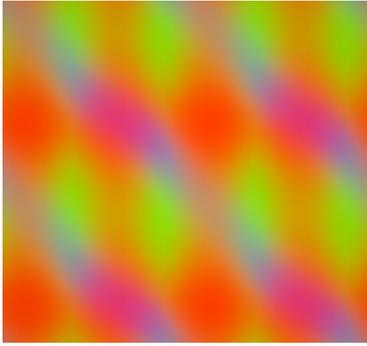
To embody error diffusion, we re-calculate M_2 after each main iteration.

The final step in the algorithm is to re-integrate the gradients to form a greyscale image. When the curl of the image is non-zero at any point, there will be no greyscale image whose gradient exactly matches that of the colour image. We therefore compute the image whose gradient is closest to the gradient of the original in the least-squares sense. The solution to this problem is given by the solution to a Poisson's equation, which we solve using the Frankot-Chellappa method.¹⁵ With homogeneous Neumann boundary conditions, this leaves a free additive constant, which we deal with by simply shifting and scaling the image to fill the interval $[0 \ 1]$.

4. RESULTS

To illustrate the basic operation of the Minkowski-luminance approach, we artificially generate an image whose L_2 norm is uniform (shown in Figure 3(a)). The sign-assignment using L_2 -type luminance gives a large average curl (i.e., integrability error) value, and the resulting greyscale image embodies very little structure of the colour original (Figure 3(b)). L_1 and L_5 (Figure 3) based sign assignments (Figures 3(c) and (d) respectively) both significantly reduce the curl and retain much more image structure from the original.

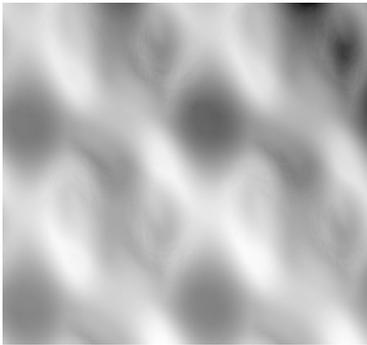
Likewise we can generate an artificial image to demonstrate the Markov relaxation (MR) approach. The image in Figure 4 (a) shows a red-green step-edge, and we imagine that the sign-assignment is unknown along the edge. Figure 4(b) shows the resulting greyscale when signs are chosen randomly along the edge; there are clear visual artefacts. The Markov relaxation approach outlined here identifies the problem pixels along the edge and visits them randomly, changing the sign to minimise the integrability error. After several iterations the mean curl value reaches zero, and the image in 4(c) results.



(a) Original image



(b) L_2 sign assignment, mean curl = 0.0174



(c) L_1 sign assignment, mean curl = 0.0059

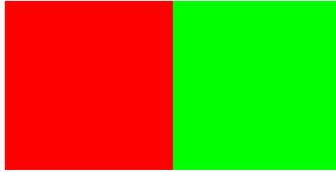


(d) L_5 sign assignment, mean curl = 0.0030

Figure 3. Effect of a Minkowski-norm sign assignment on non-integrability correction.

In Figure 5 we show Monet’s “Impression: sunrise” and show the performance of the two-stage approach. In the first image Figure 5(a) the L_1 -norm is used for sign-assignment, and the sun does not stand out from the background. In Figure 5(b) we use the L_5 norm: this reduces the overall curl value in the image, and results in a much more visible sun. We recall here that the *magnitude of the contrast is identical* for both vector fields that give rise to these images: *only the signs are different*. Finally in Figure 5(c) we show the effect of the MR on the L_5 -sign image. Although the effect is not strongly visible, Figure 5(d) shows that the MR has altered pixels around the sun, with a resulting reduction in mean curl value across the image.

The synthetic “colourwheel” image in Figure 6 is created with uniform L_1 gradient across the image, leading in an L_1 -based sign-assignment to a completely incorrect reconstruction, without our gradient-sign correction. The original image is shown in Figure 6(a); incorrect re-integration is shown in Figure 6(b), with our reconstruction shown in Figure 6(c). For comparison, the less expressive L_2 luminance image is shown in Figure 6(d): this clearly does not well represent the colour original since it blends differing colours, whereas the proposed method does a much better job. The image in Figure 7 again shows the preferability of the proposed approach.



(a) Original image



(b) Greyscale without MR



(c) Greyscale with MR

Figure 4. A simple image and two greyscale reproductions formed by re-integrating the gradient field.



(a) L_1 sign-assignment, mean curl = 0.00203



(b) L_5 sign-assignment, mean curl = 0.00195



(c) L_5 sign-assignment, post MR, mean curl = 0.00191

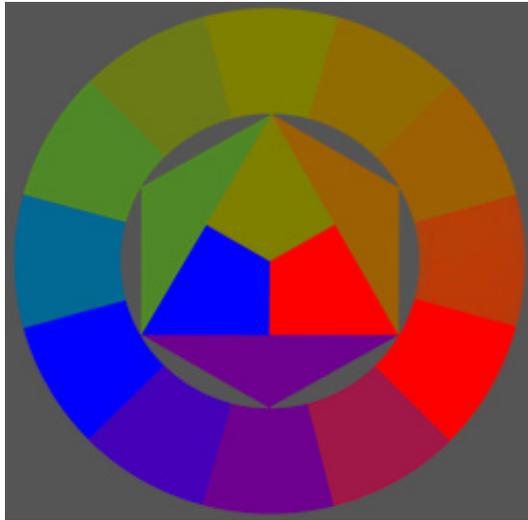


(d) Difference between pre- and post-MR

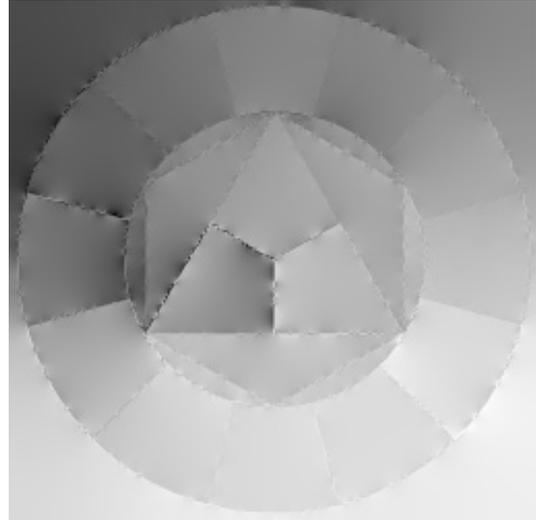
Figure 5. Both a Minkowski-norm sign assignment and a Markov relaxation makes a substantive difference to the greyscale re-integrated result.

5. CONCLUSIONS

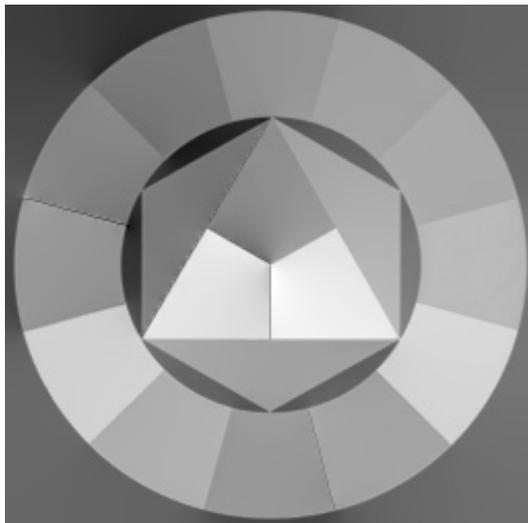
In this paper we address the sign ambiguity inherent in an eigenvector based approach to generating a representative greyscale image from a colour original. In the first place, we make use of a Minkowski-norm based decision process for assigning the vector sign for a putative grey gradient such that its direction agrees with that of increasing L_p luminance. This stage has the tendency to remove many ambiguities stemming from equiluminant edges commonly occurring in images. In a second step, we carry out a Markov relaxation method whereby pixels with problematic integrability artefacts are re-assigned a gradient vector sense so that integrability is



(a) Original image



(b) Greyscale without L_p or MR



(c) Greyscale with MR



(d) L_2 image

Figure 6. MR decreases integrability errors: synthetic image with uniform L_1 -luminance.

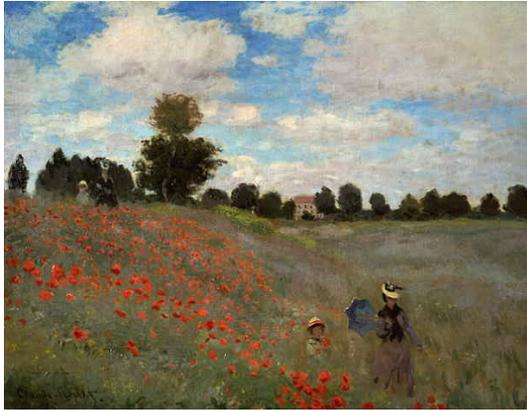
maximised locally for that pixel and its proximate neighbours. Visiting these pixels randomly and iterating several times tends to substantially reduce integrability errors. Results are shown to be considerably preferable grey representatives of the input colour images.

6. ACKNOWLEDGEMENTS

This work was supported by the Natural Sciences and Research Council of Canada and by EPSRC grants EP/E012248/1 and EP/E12159/1.

REFERENCES

1. A. A. Gooch, S. C. Olsen, J. Tumblin, and B. Gooch, “Color2gray: salience-preserving color removal,” in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pp. 634–639, ACM Press, 2005.



(a) Original image



(b) Greyscale without L_p or MR



(c) Greyscale with MR



(d) L_2 image

Figure 7. Again, performance of L_p norm and MR clearly decreases integrability errors, and also surpasses the luminance.

2. T. Hansen and K. R. Gegenfurtner, "Co-occurrence of color and luminance edges in natural scenes.," in *Bild- und Signalverarbeitung Ilmenau e. V. (Ed.), 12. Workshop Farbbildverarbeitung.*, (Ilmenau, Germany), October 2006.
3. C. L. Novak and S. Shafer, "Color edge detection.," in *Proceedings DARPA Image Understanding Workshop, Vol. 1.*, pp. 35–37, (Los Angeles, CA, USA), February 1987.
4. R. Bala and R. Eschbach, "Spatial color-to-grayscale transform preserving chrominance edge information," in *IS&T/SID 12th Color Imaging Conference*, pp. 82–86, IS&T/SID, (Scottsdale, Arizona), 2004.
5. M. Grundland and N. A. Dodgson, "The decolorize algorithm for contrast enhancing, color to grayscale conversion," Tech. Rep. UCAM-CL-TR-649, Computer Laboratory, University of Cambridge, October 2005.
6. K. Rasche, R. Geist, and J. Westall, "Detail preserving reproduction of color images for monochromats and dichromats," *IEEE Computer Graphics and Applications* **25**, pp. 22–30, May 2005.
7. L. Neumann, M. Cadik, and A. Nemcsics, "An efficient perception-based adaptive color to gray transformation," in *Proceedings of Computational Aesthetics 2007*, pp. 73–80, Eurographics Association, (Banff, Canada), 2007.
8. D. Connah, G. D. Finlayson, and M. Bloj, "Seeing beyond luminance: A psychophysical comparison of techniques for converting colour images to greyscale," in *IS&T/SID 15th Color Imaging Conference*, pp. 336–340, IS&T/SID, (Albuquerque, New Mexico), 2007.
9. D. A. Socolinsky and L. B. Wolff, "Multispectral image visualization through first-order fusion.," *IEEE Trans. on Image Proc.* **11**(8), pp. 923–931, 2002.

10. D. A. Socolinsky, *A variational approach to image fusion*. PhD thesis, Johns Hopkins University, 2000.
11. S. Di Zenzo, "A note on the gradient of a multi-image," *Computer Vision, Graphics, and Image Processing* **33**, pp. 116–125, January 1986.
12. A. Agrawal, R. Chellappa, and R. Raskar, "R.: An algebraic approach to surface reconstruction from gradient fields," in *In: Proc. Int. Conf. Computer Vision*, pp. 174–181, 2005.
13. G. D. Finlayson and E. Trezzi, "Shades of gray and colour constancy," in *Color Imaging Conference*, pp. 37–41, 2004.
14. A. Alsam and M. S. Drew, "Fast colour2grey," in *IS&T/SID 16th Color Imaging Conference*, pp. 342–345, IS&T/SID, (Portland, Oregon), 2008.
15. R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.* **10**(4), pp. 439–451, 1988.