A Unified Approach to Colour2grey and Image Enhancement Through Gradient Field Integration

Graham D. Finlayson, Roberto Montagna and David Connah; University of East Anglia; Norwich, UK

Abstract

The standard approach for the conversion of colour images to greyscale is to remove their chromatic content and keep only the brightness information. However, with this method the detail at equiluminant edges disappears. In this paper, we propose a colour-to-greyscale transformation that preserves details and, at the same time, enhances the output. We consider Socolinsky and Wolff's technique that codes the contrast information of the three colour channels in a single gradient field. Our contribution is to show how retinex-type algorithms produce a greyscale image by integrating such a gradient. Our approach also addresses the question of integrability where the gradient field is non-integrable (causing the greyscale to have smears and bending artefacts). We show how by diminishing the saturation of the original input we not only reduce the integrability problem but, with enhancement, keep the perceived colour contrast in a greyscale reproduction.

Introduction

It is more and more common to capture images using colour digital cameras, and to reproduce them on colour monitors or through colour inkjet printers. However, it is also common that they are reproduced on black and white devices, such as printers, photocopiers, fax machines; moreover, often newspapers print their pages in black and white for cost reasons.

In order to perform the necessary conversion, the colour images are deprived of their chromatic content, thus what remains is only the brightness information. Although this approach often produces acceptable results, it is desirable to have a better greyscale reproduction, that preserves more detail in those equiluminant areas where some colour contrast is present. Not only this would allow a more pleasant reproduction through black and white devices, it might also have applications as an aid for colour blind people.

Although researchers have recently devised different methods to tackle this problem [8, 9, 13, 15], there is a common basis in what has been done. A lot of effort was put in defining a measure for colour differences that could then be mapped into greyscale differences in the final image. The most general and mathematically well-founded contribution in this sense is the work of Socolinsky and Wolff [16]: they code the local colour contrast by defining a gradient field, that can be integrated to obtain a greyscale image. This is the point where we start in this paper: first of all we will discuss Socolinsky and Wolff's definition of the gradient field for multichannel images. Then we will briefly overview the principles behind retinex algorithms, and we will see how to use them to integrate a gradient field, thus obtaining a greyscale image. After that we will see how adjusting the saturation of the original colour image affects the integrability of the gradient field, and finally we will show some results of our colour-to-greyscale method.

Colour contrast

Di Zenzo [3] and, some years later, Socolinsky and Wolff [16] introduced an elegant mathematical formulation of the contrast for multichannel images. In their work, Socolinsky and Wolff define a single gradient field from an image of more channels. A colour image can be represented as a function $f\colon \Omega \subset \mathbb{R}^2 \to \mathbb{R}^3$, mapping a region of the 2D plane $\mathbb{R}^2$ to a colour space in $\mathbb{R}^3$. The gradient of the image at a point (x_0, y_0) is given by the Jacobian matrix

$$J = \begin{bmatrix} \frac{\partial f_1}{\partial x} & \frac{\partial f_1}{\partial y} \\ \frac{\partial f_2}{\partial x} & \frac{\partial f_2}{\partial y} \\ \frac{\partial f_3}{\partial x} & \frac{\partial f_3}{\partial y} \end{bmatrix}, \tag{1}$$

where the function $f_i(x,y)$ represents the ith colour channel. The matrix J can be used to compute the gradient in an arbitrary direction $d = [\cos\theta, \sin\theta]^T$ as Jd, whose magnitude is given by

$$m = d^T J^T J d. \tag{2}$$

The matrix $J^T J$ is called the *structure tensor*, and can be used to find the direction d where the magnitude m of the gradient is the largest. Such direction, as outlined in [16], is given by the eigenvector of the matrix $J^T J$ corresponding to the largest eigenvalue, whose square root is then its magnitude. Unfortunately, this method does not give any information about the sign of the gradient: this problem is serious, because it is the sign that determines whether the gradient is going from a dark to a light area, or the other way round. Thus, we can take the sign from the luminance image, or the ℓ^2-norm luminance; in our work we use the ℓ^∞-norm, that is the $\max(R,G,B)$ image, for the sign assignment. Again, in all these cases there could be equiluminant areas, even in a ℓ^2 or ℓ^∞ sense. Drew et al. [4] investigated in detail this problem, trying to address the ordering with different ℓ^p-norms and with a Markov relaxation technique, but the discussion of these issues goes beyond the purpose of this paper.

Integration of the gradient field

The main contribution of Socolinsky and Wolff's work, compared to Di Zenzo's, is that they suggest integrating this gradient field in order to get a greyscale image that can be visualised. Their approach, which is more general than the colour-to-greyscale problem, was devised for applications such as displaying satellite images, that are taken with more than three channels (not necessarily in visible light). The main idea is to see the integration problem as the solution of a Poisson equation: if the gradient obtained from a colour image I is $G = (G_x, G_y)$, the greyscale image L is the solution to the equation

$$\nabla^2 L = \frac{\partial G_x}{\partial x} + \frac{\partial G_y}{\partial y}, \tag{3}$$

where ∇^2 is the Laplacian operator. However, very often the gradient field G is not integrable, that is to say that there is not such an image L whose gradient is G. Socolinsky and Wolff suggest integrating the gradient field with an iterative relaxation method, in order to obtain an approximation to the solution in a least-squares sense. Others have approached the same problem, although for different applications, with other solutions: for example, Fattal et al. [5] suggest solving the Poisson equation from a non-integrable gradient field using multigrid methods (see e.g. [14]), while Borenstein [1] uses a method based on the Fourier transform.

Retinex algorithms

Retinex algorithms, to some extent, consist of the separation of illumination from reflectance. Land and McCann [11] suggest that the perception of colour is mainly due to the latter, rather than the former. In their paper [11], they consider a Mondrian scene and a piece-wise linear curve across different patches, thus for the moment we can consider a one-dimensional image $i(x) = r(x)l(x)$, where $r(x)$ represents the reflectance and $l(x)$ the illumination. Land and McCann assume that across an image the reflectance varies sharply at the border between objects, whereas the illumination changes smoothly. Therefore, in order to separate the two effects at a certain edge, the following ratio between areas x_2 and x_1 should be considered:

$$\frac{i(x_2)}{i(x_1)} = \frac{r(x_2)l(x_2)}{r(x_1)l(x_1)} \approx \frac{r(x_2)}{r(x_1)}. \tag{4}$$

The last equality in equation (4) follows from the fact that $l(x_2) \approx l(x_1)$. However, the human eye cannot compare directly two areas of the scene that lie very far away from each other. Thus, Land and McCann propose to multiply the ratios of all the boundaries between the two areas:

$$\frac{i(x_n)}{i(x_1)} = \frac{i(x_2)}{i(x_1)} \cdot \frac{i(x_3)}{i(x_2)} \cdot \ldots \cdot \frac{i(x_n)}{i(x_{n-1})} \approx \frac{r(x_n)}{r(x_1)}. \tag{5}$$

In practice, such a sequence of products and divisions is computationally very expensive (especially in the years when this method was devised), hence it is wiser to consider the logarithm of the image in order to transform products into sums and divisions into differences, so if $I(x) = \log i(x)$, equation (5) becomes

$$\begin{aligned} \log \frac{i(x_n)}{i(x_1)} &= \log i(x_n) - \log i(x_1) = I(x_n) - I(x_1) = \\ &= \big(I(x_2) - I(x_1)\big) + \big(I(x_3) - I(x_2)\big) + \ldots + \big(I(x_n) - I(x_{n-1})\big). \end{aligned} \tag{6}$$

Borenstein [1] suggests that the local differences approximate the derivative of $I(x)$, and summing these approximates their integration.

Over the years, several variations of the original retinex formulation have been proposed. Here are the main operations in the different versions of the algorithm, apart from the ratio-and-product step just described.

- *Threshold operation.* If the purpose of retinex is to separate the contribution of the illumination from that of the reflectance, working under the assumption that the illumination changes very slowly across the scene, when a ratio is very close to 1 it is forced to assume a value of 1. In the logarithm of the image, the ratio step becomes $\delta(I(x_{k+1}) - I(x_k))$, where

$$\delta = \begin{cases} 0 \text{ if } |I(x_{k+1}) - I(x_k)| \leq \tau \\ 1 \text{ if } |I(x_{k+1}) - I(x_k)| > \tau \end{cases} \tag{7}$$

  with τ threshold.
- *Reset operation.* In order to estimate more accurately the reflectance of a certain area, this has to be compared to a very bright area within the scene, that is assumed to have reflectance 1. However, the position of such an area is unknown, therefore it can happen that the sequence of ratios comes across a brighter area than its starting point. If this is the case, the overall product computed so far will have value greater than 1 (or, in logarithm, the overall sum will be greater than 0), thus a reset is triggered and the sequence of ratios restarts assuming reflectance 1 in the current position.
- *Average operation.* When estimating the reflectance of an area i_n, retinex starts from a point i_1 and performs ratios and products along some curve that leads to i_n. This idea can be extended by considering several curves starting from different areas: in this case, it is necessary to average (in logarithm space) the different reflectances obtained from the different curves. Frankle and McCann [6] introduced a different concept of average: every time a certain area is reached, the computed reflectance is averaged with the reflectance computed on the previous iteration.

So far, for the sake of simplicity we have referred to path based retinex algorithms, that in some sense are one-dimensional. In the literature there are several examples of two-dimensional retinex. For example, Frankle and McCann [6] consider differences between pixels at different and increasingly small distances, separating the horizontal and vertical directions. As the same pixels are taken into consideration more than once, Frankle and McCann introduced the aforementioned averaging step. McCann [12], using a slightly different approach, considers on each pixel the ratios with the eight neighbouring pixels, and the operation is repeated in several scales of the image. In the limit (infinite steps), Frankle and McCann's algorithm returns the original image (see Brainard and Wandell [2]). For a smaller number of steps per scale, it tends to enhance the detail in the image.

Integrating a gradient field using retinex

In this paper, we propose to integrate with a retinex algorithm the gradient field $G = (G_x, G_y)$ obtained with Socolinsky and Wolff's method. We base our work on Frankle and McCann's retinex and its implementation by Funt et al. [7]. Our technique consists, first of all, in computing the gradient of an input colour image according to Socolinsky and Wolff's method. As retinex works with logarithms, this operation is also performed on the logarithm of the input image. At this point, we have a gradient field $G = (G_x, G_y)$ that substitutes the retinex ratio step. In

Frankle and McCann's retinex it is straightforward to operate this substitution, because it separates the two directions in its ratio-product computation. As mentioned before, this algorithm considers the ratios of pixels at different distances from each other; however, our gradient is computed between neighbouring pixels. In order to overcome this limitation, in our implementation we approximate long-distance x or y derivatives by summing local derivatives along paths in the x and y directions respectively. For example, if G_x is in a form similar to the following:

$$G_x = [I_2 - I_1, I_3 - I_2, I_4 - I_3, \ldots, I_n - I_{n-1}]. \quad (8)$$

and we want to know the difference $I_4 - I_1$, we just need to add the first three coefficients in equation (8). This statement is true if G is integrable; but as we know, this is not always the case for gradients obtained with Socolinsky and Wolff's method.

Once the ratio step has been substituted with the described procedure, the rest of the algorithm follows Frankle and McCann's retinex: starting from a white image, we proceed with the product step between different parts of the gradient, with the reset step and finally we average with the previously computed product.

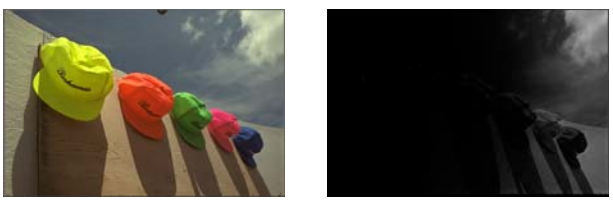

(a) Colour image (b) Retinex integrated grescale

Figure 1. *A colour image (left) and its greyscale (right) version obtained integrating the Socolinsky and Wolff gradient with Frankle and McCann's retinex. In presence of a large magnitude in the gradient field, retinex produces a halo artefact, that in this case is so heavy that hides completely a part of the image.*

What happens if we apply this method directly to a colour image? In many cases the output image is poor. In figure 1 we show an example where we do not obtain the desired effect. Image 1(a) is taken from the Kodak true colour image suite [10], and image 1(b) is the straight integration of its Socolinsky and Wolff gradient using four iterations of Frankle and McCann's retinex. What we can see there is a huge halo artefact that expands from the first hat on the left across the whole image: this is due to a very large magnitude in the gradient field. In the next section we will discuss this problem and our approach to tackle it.

Reducing saturation to improve integrability

When taking the logarithm of an image, zero values are problematic: $\log 0$ is not defined, thus we just replace it with the logarithm of a small number (e.g. $\log(10^{-6})$) in the pixels with zero value, resulting in a relatively large negative number. In RGB, very saturated colours have a high value in one or two channels, and a very low one in the others. Thus, on boundaries involving very saturated colours, the Jacobian matrix J of equation (1) has very large coefficients, therefore the eigenvalues of $J^T J$ will be very large as well (again, when operating with logarithms). In order to tackle this issue, we simply reduce the saturation of the input image. In our implementation, we convert the image to the HSV colour space, decrease the saturation by multiplying its channel by some coefficient smaller than 1, and then switch back to RGB. As default, we reduce the saturation to 50%, because we have seen that this gives in general good results. This approach brings the input image somewhat closer to its luminance version (that would have zero saturation), yet preserves most of its chromatic information. Moreover, the slightly reduced colour contrast will be "put-back" in the enhancement step. When computing Socolinsky and Wolff's gradient of the logarithm of a desaturated image, the magnitude issue is reduced, leaving a much smoother gradient. In figure 2(b) we can see how the unwanted dark area of figure 1(b) completely disappears. The difference between the two images is only the saturation reduction of the input: the integration method is exactly the same (four iterations of Frankle and McCann's retinex).

(a) NTSC luminance image (b) Retinex integrated greyscale

Figure 2. *Left: NTSC luminance image of the image 1(a). Right: Frankle and McCann's retinex integration of Socolinsky and Wolff's gradient obtained from the same image with saturation reduced to 25%.*

In figure 2 we can see a comparison between the luminance image of 1(a) and its greyscale version obtained with the method we are proposing. We can see how the colour gradient emphasises the difference between the hats, which is completely lost in the luminance image. Moreover, image 2(b) looks enhanced, meaning that it shows the typical features of retinex processing, including unfortunately its drawbacks: just by the first two hats there are some dark halos.

Results

We have tested our method on the Kodak colour image dataset [10], and on a few test images from Gooch et al. [8] and Rasche et al. [15]. In general, our proposed method tends to preserve the colour contrast, plus it introduces some enhancement to the input images. Thus the output is very different from a luminance image processed by some retinex algorithm, as shown for example in figure 7. Note that this is equivalent to an extreme application of the saturation reduction to our method: when saturation goes to zero, the result is a luminance image.

In the presence of inconsistencies in the gradient field, which are a source of integration errors, our retinex method tends to produce halos around the non-integrable areas, as shown in figure 2(b); however, in general, images tend to be free from artefacts. In figures 3, 4 and 6 we can see some examples of greyscale images produced with our method, compared to their luminance greyscale. As mentioned before, we can see clearly the enhancement (especially in figure 3).

(a) Colour image (b) Luminance (c) Our method with 50% saturation

Figure 3. *In the luminance image, the colour contrast between the leaves and the berries completely disappears. Our method instead preserves it, the result of this being that the berries look brighter than the leaves around them. Moreover, the image looks enhanced.*

(a) Colour image (b) Luminance (c) Our method with 50% saturation

Figure 4. *Our method clearly marks the difference between the flower in the middle and the leaves around it, that completely disappears in the luminance image. The enhancement is quite clear on the leaves, where some details pops out even in comparison with the colour image.*

(a) Colour image (b) Luminance (c) Our method with 50% saturation

Figure 5. *In this example our method shows a clear difference between areas of different colours, that disappears in the luminance image. Moreover, looking at the texture of the grass, we can see the enhancement introduced by our method.*

Conclusion and future work

In this paper we have introduced a colour-to-greyscale method that exploits retinex features in order to integrate and enhance a contrast preserving gradient field. We have shown how its output images preserve detail that is lost in their luminance greyscale. Further, although our approach can introduce artefacts (such as halos) when the gradient field is non-integrable, these can be removed by reducing the saturation in the original image.

We think that our method is very promising, and as part of its further development there are several aspects requiring further attention. First, in the literature there are other retinex algorithms that we could adopt. Second, the strong improvement introduced

(a) Colour image (b) Luminance (c) Our method with 50% saturation

Figure 6. *In this image, one could argue that its luminance version does not need any enhancement. However, we can notice how our method preserves some detail in the girl's pullover, that was invisible in the luminance image.*

(a) Colour image (b) Luminance (c) Retinex on luminance (d) Our method with 50% saturation

Figure 7. *This example shows that the simple application of Frankle and McCann's retinex on a luminance image does not produce the same result as our method, that introduces a much more vivid contrast on its output image.*

by mild saturation reduction is very meaningful and worth further study. Last but not least, in this paper we only touched upon the sign of the gradient computed with Socolinsky and Wolff's method, yet a general solution to this problem (such as the one proposed in [4]) may greatly improve the quality of produced images.

Acknowledgments

This work is supported by the EPSRC grant number EP/E012248/1.

References

[1] Eran Borenstein. Lightness and brightness computations by Retinex-like algorithms. Master's thesis, Feinberg Graduate School, Weizmann Institute of Science, Rehovot, Israel, 1999.

[2] David H. Brainard and Brian A. Wandell. Analysis of the retinex theory of color vision. *J. Opt. Soc. Am.*, 3(10):1651–1661, 1986.

[3] Silvano Di Zenzo. A note on the gradient of a multi-image. *Comput. Vision Graph.*, 33(1):116–125, 1986.

[4] Mark S. Drew, David Connah, Graham D. Finlayson, and Marina Bloj. Improved colour to greyscale via integrability correction. In *Proceedings of Electronic Imaging*. IS&T and SPIE, 2009.

[5] Raanan Fattal, Dani Lischinski, and Michael Werman. Gradient Domain High Dynamic Range Compression. In *SIGGRAPH '02: ACM SIGGRAPH 2002 Papers*. ACM, 2002.

[6] Jonathan A. Frankle and John J. McCann. Method and apparatus for lightness imaging. U.S. Patent No. 4,384,336, 1983.

[7] Brian V. Funt, Florian Ciurea, and John J. McCann. Retinex in Matlab. In *Proc. 8th Color Imaging Conference*, pages 112–121. IS&T and SID, 2000.

[8] Amy A. Gooch, Sven C. Olsen, Jack Tumblin, and Bruce Gooch. Color2gray: Salience-preserving color removal. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pages 634–639. ACM, 2005.

[9] Mark Grundland and Neil A. Dodgson. The decolorize algorithm for contrast enhancing, color to grayscale conversion. Technical Report UCAM-CL-TR-649, Computer Laboratory, University of Cambridge, UK, October 2005.

[10] Kodak. Kodak Lossless True Color Image Suite. Webpage, 2004. URL `http://r0k.us/graphics/kodak/`.

[11] Edwin H. Land and John J. McCann. Lightness and retinex theory. *J. Opt. Soc. Am.*, 61(1):1–11, 1971.

[12] John J. McCann. Lessons learned from Mondrians applied to real images and color gamuts. In *Proc. 7th Color Imaging Conference*, pages 1–8. IS&T and SID, 1999.

[13] Laszlo Neumann, Martin Čadík, and Antal Nemcsics. An efficient perception-based adaptive color to gray transformation. In *Proceedings of Computational Aesthetics 2007*, pages 73–80, Banff, Canada, 2007. Eurographics Association.

[14] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes*. Cambridge University Press, 3 edition, 2007.

[15] Karl Rasche, Robert Geist, and James Westall. Detail preserving reproduction of color images for monochromats and dichromats. *IEEE Comput. Graph. Appl.*, 25(3):22–30, 2005.

[16] Diego A. Socolinsky and Lawrence B. Wolff. Multispectral image visualization through first order fusion. *IEEE T. Image Process.*, 11 (8):923–931, 2002.

Author Biography

Roberto Montagna received his BSc in computer science in 2004 and his MSc in intelligent and multimedia systems in 2007 from the University of Verona (Italy). Since 2007 he has been a PhD student at the University of East Anglia, Norwich (UK) under the supervision of Prof. Graham D. Finlayson.