

Coupled Retinex

Javier Vazquez-Corral^{1,2,*}, Graham D. Finlayson^{1,*}

¹*School of Computing Sciences, University of East Anglia; Norwich, UK*

²*Universitat Pompeu Fabra; Barcelona, Spain*

**Both authors contributed equally to this paper.*

Abstract

Retinex is a colour vision model introduced by Land more than 40 years ago. Since then, it has also been widely and successfully used for image enhancement. However, Retinex often introduces colour and halo artefacts. Artefacts are a necessary consequence of the per channel color processing and the lack of any strong control for controlling the locality of the processing (halos are very local errors).

In this paper we relate an input to the corresponding output processed retinex image by using a single shading term which is both spatially varying and smooth and a global colour shift. This coupling dramatically reduces common Retinex artefacts. Coupled Retinex is strongly preferred in preference tests.

1. Introduction

Retinex -a portmanteau word from Retina and Cortex- is a colour vision model that was defined by Land more than 40 years ago [1, 2]. In Retinex theory it is proposed that *colour sensation* is not related to the radiance values that reach the eye, but to the integrated reflectance. This integrated reflectance is, loosely, defined as the ratio at each waveband between the value of the object and the value of a white object under the same illuminant. We say ‘loosely’ because the Retinex processing operates iteratively and locally and so white may refer to a white patch or to some local per color channel maximum. The locality of the processing allows Retinex to account for slowly varying illumination. See [1] for a review of the Retinex theory.

Assuming that the Retinex processing works (a colour response is divided by the response for a local white patch) then it can be shown that Retinex ‘solves’ the colour constancy problem. The same scene viewed under different coloured lights, post retinex processing has the same integrated reflectance value. We remark that this ‘dividing by white’ is implemented in many real camera processing pipelines.

The Retinex computational framework - which was designed to account for how we see colours - can be repurposed to solve other visual tasks. If one part of the image is bright and the other dark then post Retinex processing the colours in both regions are more similar. As such Retinex can be used to tackle the dynamic range compression problem. Moreover, the Retinex processed image also looks enhanced in terms of contrast. Retinex has also been used in algorithms for shadow removal, dehazing and gamut mapping [3, 4, 5, 6, 7, 8]. For the rest paper we will extend the basic Retinex approach for image enhancement (though we note that the modifications we propose might also prove useful for the other applications listed).

The computational structure of the foundational Frankle-McCann - the original computational algorithm and the one we

review here - is particular. There are three important concepts (which we will recapitulate in detail in the next section), the path, the threshold and the reset. The algorithm works iteratively. Values in one image region are propagated to another along a random path. As the path is swept out we remove any small variations (threshold out small differences) and we make sure any path predicted values are physically plausible (below the maximum value for the image). Retinex processing happens in all three, R, G and B channels separately. And there are many paths which iteratively contribute to the final calculated image.

The particular computational structure of Retinex leads to particular artefacts. In Figure 1 we show in the left the original image, and in the right the output from the Frankle-McCann Retinex method [9]. In this example, note the halos appearing around the fishes in the aquarium, and the unrealistic faces in the audience.

In this paper, we present an approach to remove or reduce the appearance of artefacts introduced by Retinex algorithms (including the Frankle-McCann approach and the large number of algorithms that have evolved from it). Our approach is, in effect, a simple post-processing step that relates the input and Retinex processed images to produce a second output that is more or less equally enhanced but has fewer artefacts. In detail we propose that the Retinex output can be approximated as a *local* shading and *global* colour change from the original. Additionally, in our framework we can easily add constraints such as the shading change could have bounded smoothness. By controlling both how colour is changed and the degree of locality in our processing we can often remove and always mitigate the typical Retinex image artefacts.

In section 2, we review Retinex (as we use the term in this paper). Section 3 presents our new method. Experiments are reported in section 4. The paper finished in section 5.

2. Background

Retinex was first introduced by Land to model certain aspects of human colour vision. Its main aim was to account for the color constancy phenomenon where we see objects under different coloured lights as being the same colour, i.e. the capability we have to perceive the colour of the objects stable regardless of the illumination. Formally, Retinex attempts to decouple surface reflectance from illumination [2]. Land and McCann hypothesised that we see object colours independent of the colour of the light because the Retinex process recovers the object surface reflectance (actually the per channel albedo).

Let us now give a brief overview of the original Retinex algorithm. Suppose we have a single channel intensity image I




Figure 1: Example of a Retinex artefact. Left: Original Image. Right: Image corrected by using Frankle-McCann Retinex

(in the case of a colour image we will follow the same procedure on each the three individual colour channels). Let us denote the i th pixel location as ρ_i . Now, consider the path $\rho = [y = \rho_0, \rho_1, \dots, \rho_{n-1}, \rho_n = x]$ that ends in a particular pixel of interest x . Formally, we define the lightness value for this path ρ as the sequential product of ratios along the path

$$I^{\rho}(I(x)) = \frac{I(x)}{I(y)} = \frac{I(\rho_1)}{I(y)} \cdot \frac{I(\rho_2)}{I(\rho_1)} \dots \frac{I(\rho_{n-1})}{I(\rho_{n-2})} \cdot \frac{I(x)}{I(\rho_{n-1})}. \quad (1)$$

Given a set of N paths (ρ_1, \dots, ρ_N) finishing at our pixel of interest x , the lightness estimator is computed as the mean for all these paths:

$$I(I(x)) = \frac{1}{n} \sum_{i=1}^N I^{\rho_i}(I(x)). \quad (2)$$

The actual computation is actually slightly (and necessarily) more complicated than the above explained due to two extra mechanisms: *threshold* and *reset*. Ratios close to one are thresholded out (mapped to one). Amongst other effects this means that the ratios calculated in the path computation are invariant to slowly varying illumination. As we compute values along a path we are in effect predicting the value at a location given the value at the start of the path and the intervening ratios. The maximum value in an image is 1. Which means if the value calculated along a path is larger to 1 it is reset to 1. In this case, all the sequential product up to that pixel is reset to 1. With respect to the reset only, the actual computed ratio is not $I^{\rho}(I(x)) = \frac{I(x)}{I(y)}$ but can be shown to be equal to $I^{\rho}(I(x)) = \frac{I(x)}{\max_{y \in \rho}(I(y))}$.

Finally, we make two remarks. First, that there is an art to choosing how many paths are averaged and how far paths are from one another. Second, the input image is first mapped into a logarithmic space in order to simplify computations (ratios become difference and multiplications become additions). Retinex operating in log space can be implemented to run much quicker compared with the original ratio based formalism.

3. Coupled Retinex

Although Retinex was designed to model human perception perhaps its largest application domain has been image enhancement. Indeed, the algorithm described in the last section is not

fully defined. How many paths are averaged and how do the paths relate pairs of pixels (see [10] for a discussion)? If a threshold is used what is the threshold? By tuning these parameters the resulting processed image can appear quite different. In turn, dependent on these parameters there are many image enhancements. However, in general, irrespective of how these parameters are defined the resulting enhanced images have artefacts.

There are several typical artefacts visible in Retinex processed images. First, because the Retinex computation is carried out per channel the local colour balance in the image can look wrong (compared to the original unprocessed image). One of the ideas we pursue in this paper is coupling the channels together in the computation (thereby removing this kind of problem). Second, the local contrast enhancement can introduce halo artefacts at edges. Third, there can be large scale bending effects (e.g. a uniform sky may become slowly darker or brighter across an image).

In this paper, we aim at removing or reducing the artefacts introduced by Retinex. Our first idea is to couple the computation (all channel outputs are calculated together). We do this in a *post-processing sense*. Specifically, we assume that the input to our algorithm is an image pair: the original image and the output image calculated by Retinex. Then we derive a mapping from the original to the output. Our mapping has two component parts: a global colour change and the local processing that is implemented as a local shading adjustment. The global colour change is, technically, implemented as a homography [11]. The shading manipulation is a per pixel multiplicative factor. Advantageously, in terms of how we formulate the problem we can control the smoothness (or locality) of the shading adjustment. Our new coupled Retinex either has no visible artefacts or they are significantly reduced.

Mathematically, let us represent the original image I_{or} , and the output of Retinex as I_{ret} as $N \times 3$ matrices. That is, for now, we remove the spatial structure from the images (we will have to retrofit the spatial structure later). With respect to this matrix representation, we can implement our idea of coupling (global colour transform) and shading. The coupling term will be a 3×3 matrix H post-multiplying I_{or} and the shading term is a per row scaling factor. In our optimisation we find the $N \times N$ diagonal matrix D - in linear algebra this implements per row scaling, and

Algorithm 1 ALS algorithm

- **Input:** The original image I_{or} and the Retinex output I_{ret} .
 - Define $I^0 = I_{or}$, $i = 0$
 - repeat over i
 1. $\min_{H^i} \|I_{ret} - I^i H^i\|$
 2. $\min_{D^i} \|I_{ret} - D^i I^i H^i\|$
 3. Update I . $I^{i+1} = D^i I^i$
 4. Update i . $i = i + 1$
 - Until $\|I^{i+1} - I^i\| < \epsilon$, where ϵ is a predefined tolerance value.
 - **Output:** $I^{i+1} H^i$
-

the 3×3 matrix H - the coupling term- that minimises:

$$\min_{D,H} \|I_{ret} - D I_{or} H\| \quad (3)$$

The coupled version of retinex is defined as $I_{coupled} = D I_{or} H$ the result for our approach. We note that H is an homography matrix (see [11]) (but the reader unfamiliar with this theory can think of H as a simple linear transform, as indeed it operationally is).

To solve Equation (3), we adopt an Alternative Least Squares strategy (ALS) [12]. We iteratively solve for D and H until convergence. See Algorithm 1 for a pseudo code implementation of our approach. Note that the operations computed on the ALS minimization are Least-Squares based, and therefore they are not very expensive in terms of computational cost.

We impose a further constraint for D specifically that it - well, the per pixel multiplication implied by D - should be smooth. We implement smoothness by enforcing the shading term to be represented as a linear combination of the first few terms in a DCT expansion. Our minimization becomes:

$$\min_{D^{DCT}, H} \|I_{ret} - D^{DCT} I_{or} H\| \quad (4)$$

The new smooth shading adjustment, D^{DCT} , is calculated in 3 steps. First we map the $I_{or} H$ and I_{ret} to images (with $N \times M$ pixels) $L_{or}(x,y)H$ and $L_{ret}(x,y)$ (underscoring makes clear these are RGB images, each pixel has 3 numbers) and (x,y) indexes pixel location. Now we find the shading image that minimizes

$$\begin{aligned} \min_{D(x,y)} \|D(x,y) L_{or}(x,y) H - L_{ret}(x,y)\| \\ \text{s.t. } D(x,y) = \sum_{k=1}^K \alpha_k DCT_k(x,y) \end{aligned} \quad (5)$$

where $DCT_k()$ represents the k th DCT basis image. Finally, we map the recovered shading image back to the diagonal matrix representation: $D(x,y) \rightarrow D^{DCT}$.

3.1. More general colour transforms

Our default model of colour change is a 3×3 linear transform. As proposed earlier we call this a *homography*. This is because there is a scaling ambiguity between H and D (we multiply one by a factor we can divide the other by the same factor without changing the minimization). We can, however, test a more general

‘homography’ type transform. For example, a 4×3 homography matrix (by just adding a fourth column of 1s to the input image). The main interest of this matrix is the inclusion of an *offset* term.

4. Experiments

In all our experiments we use the implementation of the Frankle-McCann Retinex model presented in [13].

Our method was run until a tolerance of $1e-8$ (the ϵ parameter in Algorithm 1) was accomplished, and 10 DCT-basis have been considered for the minimisation of D^{DCT} .

4.1. Informal visual comparison

The two main problems of the Retinex approach are that it can introduce a false global colour balance and halos at some edges. Figure 2 shows examples of both these problems. In Figure 4, left to right we show 4 images: the input, the Retinex processed output, the result of processing with a 3×3 and 4×3 homography (see 3.1 in the last section). The first row of images shows two results. First, this is an example where Retinex makes an output without visible artefacts. Second that our coupled retinex gives, visually, a very similar result.

In the next two rows there are visible artefacts. See the halos appearing around the fishes on the aquarium in the middle Retinex processed image. And in the image at the bottom there is an example of an incorrect white-balance -combined with an enhancement of noise- produced by Retinex. In both cases either couples (3×3 or 4×3) provide an output image that achieves the desired enhancement without the artefact. The 4×3 outputs appear slightly brighter (and numerically, are slightly closer to the Retinex output).

4.2. Preference

We run a pairwise forced-choice experiment. Here, an observer is presented with the original image in the center and two different methods enhancements, left and right. The observers were asked to select their preferred enhanced image. Three algorithms are tested: the Frankle-McCann retinex and the 3×3 , and 4×3 homographies (see section 3.1 for description). The position - left or right of the methods was randomly selected at each comparison. Observers were sat at approximately 80 cm from the monitor (an HP LP2480zx) in a dimly lit room (with uniform grey walls). The total number of observers was seven.

Fifteen input images were randomly selected from the databases [14, 15, 16]. Since there are 3 algorithms there are 3 pairs of algorithms. Thus, the preference experiment for one observer involves 20 (images) multiplied by 3 (algorithm pairs). That is, 60 pairwise comparisons. We convert the pairwise preferences made (the observer is forced to prefer one output or the other) using Thurstone Case 5 Law of Comparative Judgement.

Results of the experiment are presented in Figure 4. We can see that our approach with 3×3 coupling matrix is the preferred, closely followed by the 4×3 . Both of them are preferred, with statistically significance -see error bars-, over the original Retinex output.

We remark that the more parameters there are in our approximation the closer we *replicate* the Retinex original. It is from this




Figure 2: Results when considering the Frankle-McCann Retinex. From left to right: the original image, the output of the Frankle-McCann Retinex, our approach with a 3×3 matrix, and our approach with a 4×3 homography matrix.




Figure 3: Original images used for the subjective experiment

vantage point that we believed the 4×3 coupling matrix failed to produce the best results.

Part of the problem with getting too close to the Retinex output is that the images we are processing here are rendered Jpegs. Consequently, the Retinex processed images enhance the Jpeg quantisation artefacts and in making our approximation we too can model these artefacts. Indeed the 4×3 homography version has the most quantisation problems, see Figure 5 for an example. It is possible that a repeat of this experiment where quantisation is not a problem would deliver results in which our higher order model coupling matrix is preferred.

4.3. Image Metrics

Finally, we run two very well-known image quality metrics -NIQE [17] and BRISQUE [18]- in the images of the previous experiment. For both metrics, the smaller is the value the better is the image. Results for these metrics -computed as the average




Figure 4: Results of our subjective experiment by using the Thurstone Case V Law.

over all the images- are shown in Table 1. Both metrics prefer our method with the 3×3 coupling matrix, in accordance with the subjective experiment. However, for the BRISQUE metric, the Frankle-McCann Retinex obtains a better result than our 4×3 coupling matrix. This ranking is in contradiction with the results of the subjective test, making us doubtful on the adequacy of using image metrics to evaluate this type of approaches.

Other models of Retinex

In Figure 6 we show, from left to right, an original image, the output of a particular Retinex model, our approach with a 3×3




Figure 5: From left to right: Our result with a 3×3 coupling matrix, and a 4×3 coupling matrix. Please note that the quantisation artefacts are more prominent for the 4×3 case.



Figure 6: From left to right: the original image, the output of a particular Retinex model, our approach with a 3×3 matrix, and our approach with a 4×3 homography matrix. The Retinex models presented in each row are (from top to bottom): the McCann99 Retinex [19], the MultiScale Retinex [20], the LIME model [14], and the SRJE Retinex.

	Frankle-McCann	Ours 3×3	Ours 4×3
NIQE	2.86	2.67	2.74
BRISQUE	27.13	23.95	28.18

Table 1: NIQE and BRISQUE values for the outputs obtained using the 20 original images shown in Figure 3. Results are computed as the average over the images.

matrix, and our approach with a 4×3 matrix. The Retinex models presented in each row are (from top to bottom): the McCann99

Retinex [19], the MultiScale Retinex [20], the LIME model [14], and the SRJE Retinex [21].

In the case of the McCann99 Retinex -first row of the figure- the output of the Retinex model is unrealistic, both in the main building and in the sky. Our coupled approximation attenuates this problem. In the case of the MultiScale Retinex -second row of the figure- we can clearly see the excessive enhancement presented in the Retinex solution, and how this problem is clearly attenuated in our solutions. In this particular case, the result in the hill region is better for the 4×3 case, while the sky is better in the 3×3 case.

In the case of the LIME model -third row of the figure-, we can see how the result of the original model is both noisy and also presents a yellowish tint. Again, both of our approaches are able to reduce the noise presented -specially the 4×3 case-, and to remove the yellowish tint.

Similarly, in the case of the SRIE Retinex, the ceiling in the image is very noisy, and the bulb presents an excessive red hue. Our results reduce the noise in the ceiling and also slightly correct the excessive red hue of the lamp. In this case, this red hue is better corrected in the 3×3 case, as this case is more restrictive with respect to the colours of the image.

Let us finally note that the general appearance of our final results is still biased by the colours presented in the Retinex solution, and for this reason we are -for example- not able to recover a full blue sky in the first image of the last figure.

5. Conclusion

In this paper, we have presented a post-processing method to reduce artefacts present in the results of Retinex models. Our approach proposes that the Retinex output can be approximated as a *local* shading and *global* colour change from the original input. We propose to look for these terms using the Alternative Least Squares (ALS) technique. Our results show that our method is able to overcome the Retinex limitations qualitatively, but also when considering subjective tests. Further work will consist on running subjective tests for other Retinex models, and on studying which further constraints can be placed to both the shading and the colour change terms.

Acknowledgments

Prof Finlayson is grateful for the support from EPSRC grant P007910.

J. Vazquez-Corral has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement number 761544 (project HDR4EU) and under grant agreement number 780470 (project SAUCE).

References

- [1] E. H. Land, "The retinex theory of color vision," *Scientific American*, vol. 237, no. 6, pp. 108–129, 1977.
- [2] E. H. Land and J. J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Am.*, pp. 1–11, 1971.
- [3] G. D. Finlayson, S. D. Hordley, and M. Drew, "Removing shadows from images using retinex," in *10th Color Imaging Conference*, 2002.
- [4] A. Galdran, A. Alvarez-Gila, A. Bria, J. Vazquez-Corral, and M. Bertalmío, "On the duality between retinex and image dehazing," in *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 8212–8221, 2018.
- [5] J. Vazquez-Corral, S. Zamir, A. Galdran, D. Pardo, and M. Bertalmío, "Image processing applications through a variational perceptually-based color correction related to retinex," in *Electronic Imaging*, vol. 6, pp. 1–6, 2016.
- [6] J. J. McCann, "Color gamut mapping using spatial comparisons," in *Color Imaging: Device-Independent Color, Color Hardcopy, and Graphic Arts VI*, pp. 126–131, 2000.
- [7] S. W. Zamir, J. Vazquez-Corral, and M. Bertalmío, "Gamut mapping in cinematography through perceptually-based contrast modification," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, pp. 490–503, June 2014.
- [8] J. Vazquez-Corral and M. Bertalmío, "Spatial gamut mapping among non-inclusive gamuts," *Journal of Visual Communication and Image Representation*, vol. 54, pp. 204–212, 2018.
- [9] J. Frankle and J. McCann, "Method and apparatus for lightness imaging," in *US Patent 4,384,336*, 1983.
- [10] R. Montagna and G. D. Finlayson, "Constrained pseudo-brownian motion and its application to image enhancement," *J. Opt. Soc. Am. A*, vol. 28, pp. 1677–1688, Aug 2011.
- [11] G. Finlayson, H. Gong, and R. B. Fisher, "Color homography: Theory and applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, pp. 20–33, Jan 2019.
- [12] G. D. Finlayson, M. M. Darrodi, and M. Mackiewicz, "The alternating least squares technique for nonuniform intensity color correction," *Color Research & Application*, vol. 40, no. 3, pp. 232–242, 2015.
- [13] B. Funt, F. Ciurea, and J. McCann, "Retinex in matlab," in *Journal of Electronic Imaging*, pp. 112–121, 2000.
- [14] X. Guo, Y. Li, and H. Ling, "Lime: Low-light image enhancement via illumination map estimation," *IEEE Transactions on Image Processing*, vol. 26, pp. 982–993, Feb 2017.
- [15] K. Ma, K. Zeng, and Z. Wang, "Perceptual quality assessment for multi-exposure image fusion," *IEEE Transactions on Image Processing*, vol. 24, pp. 3345–3356, Nov 2015.
- [16] C. Lee, C. Lee, and C. Kim, "Contrast enhancement based on layered difference representation of 2d histograms," *IEEE Transactions on Image Processing*, vol. 22, pp. 5372–5384, Dec 2013.
- [17] A. Mittal, R. Soundararajan, and A. Bovik, "Making a completely blind image quality analyzer," vol. 20, pp. 209–212, 03 2013.
- [18] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-reference image quality assessment in the spatial domain," *IEEE Transactions on Image Processing*, vol. 21, pp. 4695–4708, 11 2012.
- [19] J. McCann, "Lessons learned from mondrians applied to real images and color gamuts," in *Seventh Color Imaging Conference*, pp. 1–8, 1999.
- [20] A. B. Petro, C. Sbert, and J.-M. Morel, "Multiscale Retinex," *Image Processing On Line*, pp. 71–88, 2014.
- [21] X. Fu, D. Zeng, Y. Huang, X. P. Zhang, and X. Ding, "A weighted variational model for simultaneous reflectance and illumination estimation," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2782–2790, June 2016.

Author Biography

Javier Vazquez-Corral received the Ph.D. degree in computer science from the Universitat Autònoma de Barcelona, Bellaterra, Spain, in 2011. He is currently a post-doctoral researcher with the Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona. His research interests are related to the use of color in image processing and computer vision problems. He is also interested in bridging the gap between color in the human brain and its use in computer-vision applications.

Professor Graham Finlayson is the Director of the Colour & Imaging Lab at the University of East Anglia. He has published over 200 conference papers (many at this conference) and over 75 journal papers. He is also the inventor of 30+ patents many of which are used in commercial products. His interests span perception, colour image processing and physics-based computer vision.