

Unifying Path and Center-Surround Retinex Algorithms

Afsaneh Karami and Graham Finlayson. University of East Anglia, Norwich, England.

Abstract

Retinex is a theory of colour vision, and it is also a well-known image enhancement algorithm. The Retinex algorithms reported in the literature are often called path-based or centre-surround. In the path-based approach, an image is processed by calculating (reintegrating along) paths in proximate image regions and averaging amongst the paths. Centre-surround algorithms convolve an image (in log units) with a large-scale centre-surround-type operator. Both types of Retinex algorithms map a high dynamic range image to a lower-range counterpart suitable for display, and both are proposed as a method to simultaneously enhance an image for preference.

In this paper, we reformulate one of the most common variants of the path-based approach and show that it can be recast as a centre-surround algorithm at multiple scales. Significantly, our new method processes images more quickly and is potentially biologically plausible. To the extent that Retinex produces pleasing images, it produces equivalent outputs. Experiments validate our method.

Introduction

The Retinex theory of colour vision was proposed by Edwin Land [15]. Although initially designed as a model to account for understanding aspects of human colour perception, the related Retinex algorithms are used in various applications, including image enhancement [23], colour correction [17], and dynamic range compression [29]. At its heart, Retinex assumes that a primary goal of colour vision is to correctly perceive the colour of objects. Concomitantly, Retinex algorithms attempt to separate the effects of illumination from reflectance so that the veridical colour of objects can be ascertained. From an evolutionary perspective, it is often argued the ability to solve tasks - e.g. to reliably separate the colour of fruit from foliage and to determine the former's ripeness [24] - that the visual system must be able to discount the illuminant colour.

Very broadly, Retinex algorithms belong to one of two groups: either they are path- or centre-surround-based. In the path-based method, the brightness ratios of adjacent points along a random path through an image region are examined to find plausible material boundaries. Illumination is assumed to vary slowly across a scene, while reflectance causes an abrupt change in signal strength when an object/material boundary is encountered, and so at material boundaries, there will be a large or small ratio (significantly different from '1', no little change). Thus, we can remove the illumination by thresholding to unity small ratios.

In path-based Retinex, adjacent thresholded ratios multiply successively so that the brightnesses in one part of an image can be compared to those in the other parts. By multiplying adjacent ratios, the Retinex computation effectively *reintegrates* the image. Because the illumination variation has been thresholded away, the reintegrated signal depends only on the objects' albedos. Where albedo is a measure proportional to reflectance, and in Retinex parlance albedo is proportional to *lightness*. Importantly, Retinex integrates along multiple paths, and the final calculated lightness is the average computed along all the paths.

See Fig. 1 as an example of path-based computation. Looking at the path start-point in patch D, there is an intensity of 118. Yet, as the path meanders upwards when it crosses into patch C the value is 100. The difference in these values is assumed to be due to a slowly varying illumination. However, in path-based Retinex any small variation in signal strength is thresholded out so in relating D to A it is only the ratios at the material boundaries that matter. As shown in the figure the product of the 3 ratios (from patch D into C, C into B and B into A) is 2.4 indicating that A is estimated to be 2.4 times *lighter* than D.

According to Marr [20] there is no requirement that a theory of vision - here colour perception - should lead to a computer algorithm that processes images in the same way as the visual system. What matters is that the theory accounts for the visual data. Importantly, there is psychophysical data that path-based Retinex can predict. But, it is clear that the human system - at least at the lowest levels - is not carrying out path-based processing. So, to the extent that Retinex is an accurate theory of colour vision, it is not implemented as a path-based computation by our visual system.

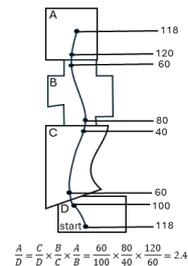


Figure 1: We calculate a path from bottom to top by multiplying ratios of luminance at material edges.

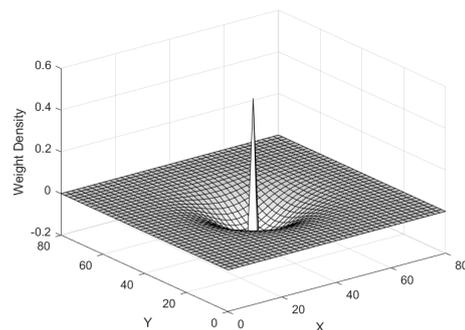


Figure 2: Convolution filter for centre-surround Retinex.

In the early stages of human visual processing, light is recorded by the cones [25], and then this signal *flows* through many intricate layers of cells connecting these sensors to the optic nerve. The receptive fields in retinal ganglion cells receive inputs from multiple cones, allowing for the transmission of contrast information [18]. These cells are organized into centre-surround receptive fields consisting of a central region and a surrounding annulus. Center-surround processing is found both in

the retina and in the visual cortex [8].

In analogy to this physiology, the centre-surround Retinex has evolved [14]. In the centre-surround Retinex, the weighted average signal in an annular region is subtracted from the brightness at the centre, where the brightness is in the log domain. In the non-log domain, the centre-surround Retinex *interprets* the image as its signal relative to (divided by) the average of the signal in the surrounding annular region. See Fig. 2 for an illustration of the centre-surround Retinex’s convolutional filter (similar to [11]). Empirically, the centre-surround Retinex mitigates the effects of illumination and can also describe some psychophysical data [14].

In this paper, we link centre-surround and path-based Retinexes. First, we can recast the commonly used McCann99 [22] path-based Retinex in the form of a coarse to fine resolution centre-surround Retinex (a type of multigrid algorithm). Second, we show that this recasting of path-based Retinex in a centre-surround framework can be interpreted explicitly as Jacobi-iteration [21] where we interpret the centre surround operator as a Laplacian - second derivative - over a variety of scales. *The extent to which Retinex mitigates for the effects of illumination and compresses the dynamic range in an image is related to the extent to which the image is partially reintegrated from its Laplacian.* This, we believe, is a wholly new and novel insight.

Of practical import, we show that our centre-surround method results in quicker processing compared to McCann99 and, to the extent that Retinex produces pleasing images, produces equivalent outputs. Qualitatively, our new centre-surround Retinex produces fewer artefacts than McCann99.

Background

In [13] a grey-scale Mondrian (a piece-wise uniform patchwork of overlapping square patches) was shown to observers where there was a very strong brightness gradient across the image. The gradient was so strong that the brightness measured from a white surface was less than for a dark patch (because there was so much more light incident on the dark patch). Remarkably, observers could correctly classify the correct surface albedos (or lightnesses) despite the large illumination gradient. To explain this and other psychophysical data, Land [15] proposed a path-based-computation. Along the path small differences due to an illumination gradient are thresholded out and only ratios - calculated at material boundaries - are propagated. Fig. 1 illustrates a path-based computation. Here patch D is related to A by multiplying only the ratios between patches. Moreover, when calculating the lightness (albedo) at A, by choosing a path that starts from D, we assume that the start point (D) has a maximal albedo of 100%. If the calculated lightness is more than 1 (more than 100% albedo) then this calculated lightness is *reset* to 1. In terms of our example, we must reset the calculated 2.4 to 1 because an albedo greater than one makes no physical sense. Note, however, if we run the path in the opposite direction then the computed value is 1/2.4. In general, the output of Retinex processing is the average of multiple random paths with random start point computations.

Let us denote an image brightness at location (x,y) as $I^l(x,y)$. According to Retinex theory the ratio $R^l(x,y)$ calculated with a neighbour, e.g. $I^l(x+1,y)$ (one pixel apart) is equal to

$$R^l(x,y) = \frac{I^l(x,y)}{I^l(x+1,y)} \quad (1)$$

Or, in the log-domain (where the absence of l indicates log-

units):

$$R(x,y) = I(x,y) - I(x+1,y) \quad (2)$$

According to the original Retinex theory, small ratios are thresholded to *no change* (though, in McCann99 - that we present below - this threshold is 0, no thresholding). In the log-domain (remembering $\log(1) = 0$)

$$R(x,y) = \text{thresh}(R(x,y)) \begin{cases} 0, & |R(x,y)| < th \\ R(x,y), & \text{otherwise} \end{cases} \quad (3)$$

Where $|\cdot|$ denotes the absolute value operator. Denoting $R_{k,i}$ as the i th ratio along a k th path ending at $I(x,y)$ the calculated ‘clipped to 0 output’, $O_k(x,y)$ is written as

$$O_k(x,y) = \min\left(\sum_i R_{k,i}, 0\right) \quad (4)$$

Let the shorthand $\text{randomPath}(I;x,y)$ denote a random path for image $I()$ that ends at location (x,y) (randomPath is an algorithm rather than a function). At this point, it is useful to reflect on what information is propagated. Indeed, we note, trivially, that $\text{randomPath}(I;x,y) + 0 = \text{randomPath}(I;x,y)$. Here 0 denotes the putative value at the start of the path. However, if we had some prior information about the image, we’d like Retinex to propagate $\text{randomPath}(I;x,y) + \text{prior}(x_k,y_k)$ (where (x_k,y_k) is the location in the image where the k th patch starts). The final output log image is the average calculated along N paths:

$$O(x,y) = \frac{\sum_{k=1}^N \min(\text{randomPath}(I;x,y), 0)}{N} \quad (5)$$

The McCann99 Retinex[22] diverges from a classical path-based formulation in three important respects. First, the path considered is one pixel long (we only look at ratios of neighbours) and, second, we always propagate a prior. Further, averaging is computed after each individual ratio propagation. The key steps in the McCann99 Retinex computation are presented in Algorithm 1.

Algorithm 1 McCann99 Retinex (computation for Level L)

- 1: Initialise $R_{u,v}^L(x,y) = I^L(x,y) - I^L(x+u,y+v)$ where $u,v \in \{-1,0,1\}$
 - 2: Initialise $O^L(x,y) = \uparrow O^{L-1}(x,y)$
 - 3: **for** $i=1$:iterations **do**
 - 4: **for** $u \in \{-1,0,1\}$ **do**
 - 5: **for** $v \in \{-1,0,1\}$ **do**
 - 6: **if** $u \neq 0$ or $v \neq 0$
 - 7: $O^L(x,y) = \frac{1}{2}(O^L(x,y) + \min(R_{u,v}^L(x,y) + O^L(x+u,y+v), 0))$
 - 8: **end**
 - 9: **end for**
 - 10: **end for**
 - 11: **end for**
-

The algorithm is hierarchical with the superscript L denoting *level*. The full resolution input image is level F and is denoted $I^F(x,y)$. The image at level $L-1$ is a half resolution downsample variant at level L . We denote this as $I^{L-1}(x,y) = \downarrow I^L(x,y)$. While the input images are downsampled the outputs we calculate are upsampled to level. The output at level L is initialised as the output from level $L-1$ upsampled by a factor of 2, denoted: $O(x,y)^L = \uparrow O^{L-1}(x,y)$. Notice that the output image that is being calculated *overwrites* itself but that the ratios $R_{u,v}^L(x,y)$ are

calculated from the input image $I^L(x,y)$ and are inviolate in each resolution. Algorithm 1 is computed per level from lowest to highest resolution ($L = 1$ to $L = F$). The coarsest level, typically a single pixel, is initialised to 0, $O^0(x,y) = 0$. The core of the path-based computation is articulated in step 7 (a ratio, here a log-difference, is used to propagate the current value at a neighbouring pixel and this is reset so it is less than 0). The result is then averaged with the current estimated output image value. The *iterations* parameter is important and is tuned for preference (iterations are typically in the range 2 to 10).

Let's finish this section with the centre-surround Retinex. Here we have a single large convolutional filter $f(x,y)$ (illustrated in Fig. 2). The centre-surround Retinex Algorithm 2, is summarised in a single step

Algorithm 2 Land centre-surround Retinex

1: $O(x,y) = I(x,y) \otimes f(x,y)$

As a final comment, we note that the path-based (McCann99 [22]) and centre-surround (Land [14]) return log images that can and should be adjusted by an additive constant (in effect, to model a constant of integration). Typically, images are adjusted so that the brightest pixel value is 0 in log units.

Method

Let us begin by denoting the convolution of $I(x,y)$ with the filter f as $I(x,y) \otimes f$. If we ignore the reset step in step 7 of the McCann99 Retinex, we can rewrite this step as

$$7: O^L(x,y) = R_{u,v}^L(x,y) + O^L(x,y) \otimes f_{uv}$$

where f_{uv} are illustrated in Fig. 3a.

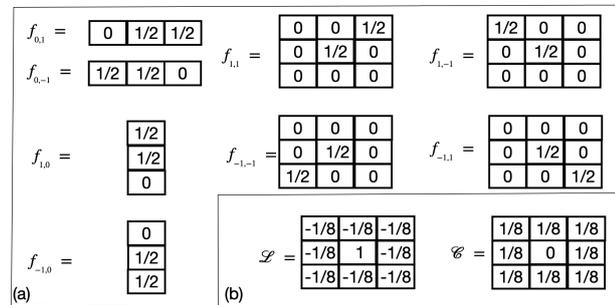


Figure 3: In (a) there are 8 filters which when convolved with an image will calculate the difference of a central pixel with each of its 8 neighbours. Panel(b), left shows a Laplacian filter and right an annular averaging filter

That we have 8 ‘directional’ averaging filters does not immediately help us view the McCann99 path-based Retinex as a center-surround computation. Rather, we choose to substitute these directional averaging with a single annulus-type operator, see filter \mathcal{C} in Figure 3b.

Rather than computing ratios (log-differences) we compute a rotationally symmetric difference. The simplest rotationally symmetric differential operator is the Laplacian [2] and it is implemented by convolving with a simple 3×3 filter \mathcal{L} (shown left of Fig. 3b): $R(x,y) = I(x,y) \otimes \mathcal{L}$. In log units, we can interpret this operation as calculating the difference between a pixel and the average of its immediate surroundings (i.e. it is a center-surround operator). We note this is one of several discrete Laplacian operators [30] (we choose the one shown mainly to illustrate the similarity of the approach we develop to McCann99, but, any typical Laplacian filter might be used).

For now, forgetting about the level L , the reset, and setting $O_0(x,y) = 0$ we define two iteration equations (6) and (7)

$$O_k(x,y) = R(x,y) + O_{k-1}(x,y) \otimes \mathcal{C} \quad (6)$$

$$O_k(x,y) = (1 - \omega)O_{k-1}(x,y) + \omega(R(x,y) + O_{k-1}(x,y) \otimes \mathcal{C}) \quad (7)$$

Where, in Equation (7) $\omega \in (0, 1]$. Here the filter \mathcal{C} (right Figure 3b) plays the same role as the directional averages in 3a). The reader familiar with algorithms in numerical analysis will notice that Eq. 6 and 7 are, respectively, a Jacobi [21] and weighted Jacobi [3] iteration. In the limit for large k then $O_k(x,y) \equiv I(x,y)$. As for McCann99 the $I(x,y)$ is recovered up to an unknown scalar (which is why we write \equiv rather than $=$). So, we must add a scalar so that the reintegrated image has a maximum of 0 (in log units). Regarding the variable ω in Equation 7, in weighted Jacobi iterations, it is sometimes necessary to set to be less than one to ensure convergence [3]. Clearly, if $\omega = \frac{1}{2}$ then the iteration in Equation 7 looks similar to step 7 in the McCann99 Retinex. Here, convergence is guaranteed for $\omega = 1$ We can now write our centre-surround version of *path-based* Retinex, see Algorithm 3. Here, as in step 7 for McCann99, we overwrite $O^L(x,y)$ and so drop the iteration subscripts i and $i-1$.

Algorithm 3 Our Centre-Surround Retinex (computation for Level L)

- 1: Initialise $R^L(x,y) = I^L(x,y) \otimes \mathcal{L}$
 - 2: Initialise $O^L(x,y) = \uparrow O^{L-1}(x,y)$
 - 3: **for** $i=1$:iterations **do**
 - 4: $O^L(x,y) = (1 - w)O^L(x,y) + w(\min(R^L(x,y) + O^L(x,y) \otimes \mathcal{C}, 0))$
 - 5: **end for**
-

The only difference between this algorithm and the classical Jacobi iteration is the *reset* (clip to 0) step (the *min* function in step 4). We note this is a sort of negative rectification and is used routinely in convolutional neural network [4] (and is proposed to have relevance to human visual physiology). Again we prove elsewhere that this version of Retinex - with reset - also converges to the original image in the limit.

As we will see in the Experiments section, our new centre-surround-path-based Retinex runs quicker than McCann99 and produces broadly equivalent results but with fewer artefacts. Moreover, and this we believe to be our most significant contribution, by recasting McCann99 as the problem of reintegrating a Laplacian field we understand the mathematical premise underpinning the algorithm. Retinex is a process for partially reintegrating an image from differential information (here Laplacian). When authors [7] have asked how many iterations should be used (step 3 in the McCann99 algorithm), they have - the perspective of our center-surround reworking of the algorithm - been asking about a Jacobi iteration and its rate of convergence. The implication of this observation is something that we are investigating.

Finally, we note that the Laplacian operator was also part of the early lightness algorithms (related to Retinex) with, for example, Horn [9] convolving an image with a Laplacian then thresholding to 0 image locations where the Laplacian was small and then reintegrating (by a Jacobi iteration). It was argued that this simple computation could predict some psychophysical results such as the Craik-Cornsweet Illusion [9]. However, in these Lightness algorithms, the Jacobi iteration was iterated until full convergence.

Experiments

The dataset of Bloch [1] comprises 11 high dynamic range, (HDR) images used for our experiments. We run the McCann99 algorithm to generate output images which look pleasing (where pleasing is entirely the subjective view of the authors). Two of the images - 'NaturalMirrorBall' and 'SantaMonicaSunset' - are shown in Fig. 4. The first row shows the original unprocessed pictures. The second row depicts rendered images rendered with McCann99. As a point of detail to make these outputs we (i) calculated the mean brightness of an RGB image, (ii) executed Retinex on the log mean brightness, (iii) exponentiated the result, and (iv) scaled so the maximum was 1 and the original HDR colour image was multiplied by the ratio of new (Retinex) over old brightnesses. In this way, we tone map only the brightness channel and preserve chromaticity with the original. Finally, (v) we applied an sRGB [27] gamma.

The first column of Table 1 shows the name of the image. The second column records the number of iterations ('its') used per image in McCann99. In brackets, we show a time in seconds to achieve processed images. The software used (Matlab) and the particular computer configuration are important issues. But, we will use the timings to (with caveats) compare the performance with our new center-surround Retinex (CS-Retinex). The number of iterations required by our new center-surround (CS) Retinex, Algorithm 3, together with timings are shown in the 3rd column of Table 1. Finally, per image, the 4th column records the mean Delta E error, CIEDE2000 [19], between images rendered by McCann99 and the new Center-Surround Retinex, similar to [10], [6], and [16]. For many image pairs a mean error of less than 3 Delta E *means* the images look the same by the human eye [26, 12]. In row 3 of Fig. 4, we show 'zoomed-

Table 1: Performance of McCann99 vs CS Retinex

HDR Image Name	McCann99 #its (time/s)	CS Retinex #its (time/s)	ΔE
CarWall	2 (5.18)	10 (0.28)	1.8
CoffeeShop	2 (5.09)	16 (0.39)	2.9
Egyptian	2 (5.12)	15 (0.38)	1.8
Engines	4 (5.55)	20 (0.47)	2.1
FatCloud	2 (5.53)	10 (0.26)	2.1
KitchenWindow	1 (4.90)	10 (0.27)	1.6
MansChinese	2 (5.09)	10 (0.27)	1.9
NaturalMirrorBall	2 (5.20)	12 (0.31)	1.6
PopcornCounter	2 (5.10)	12 (0.30)	1.6
WalkOfFame	3 (5.31)	18 (0.43)	1.8
SantaMonicaSunset	1 (4.91)	8 (0.23)	2.1

in' crops of areas in the image where there are spatial artefacts with McCann99 Retinex. Row 4 of Fig. 4 shows the result of our Center-Surround Retinex, and row 5 depicts the same previous crop areas created by our Retinex, which have reduced though not entirely removed artefacts. Importantly, [28] demonstrates that a Retinex algorithm that makes artefacts can still be used in an enhancement framework that delivers artefact free outputs.

Regarding, the timings, it is apparent our new center-surround Retinex is generally about 10 times faster despite running for more iterations. We note that in each iteration in McCann99, there are 8 separate convolutions as opposed to the single convolution in our algorithm. Significantly, in McCann99, there are 8 ratio images (per level) that need to be computed (requiring indexing and storage). In contrast, there is a single (per scale) ratio image in our new centre-surround Retinex. Also at play is the fewer numbers of resets required in the new approach. Likely, both approaches could be further optimised for imple-

mentation but, we estimate - assuming the evidence in Table 1 is representative - that our new center-surround Retinex will be at minimum twice as fast.

As a final comment, the Bloch [1] dataset is clearly small (and not representative of HDR images). However, our initial investigation of larger HDR sets, e.g. the RIT Photographic survey [5] shows a similar data trend.

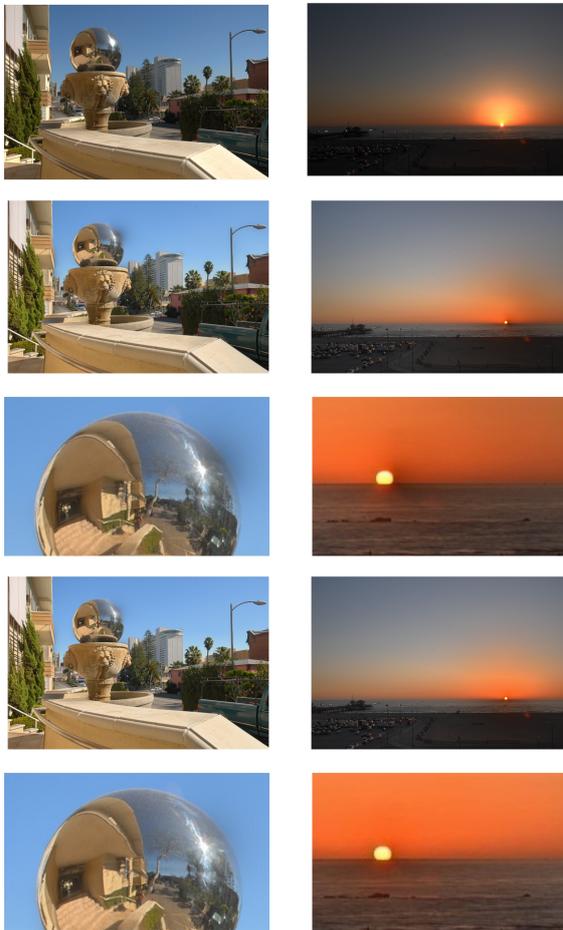


Figure 4: Row 1 is original images, rows 2 and 3 output of McCann99 Retinex and detail of artefact, rows 4 and 5 render images from our Retinex with the same crop areas.

Conclusion

There are two broad classes of Retinex algorithms: those that are based on path-based-processing and, the more biologically plausible, center-surround algorithms. In this paper, we showed that the path-based McCann99 [22] algorithm can be recast as a center-surround algorithm. Thus, this paper unifies two disparate strands of Retinex theory. Our new center-surround algorithm is also simpler, faster, and produces similar images to the prior path-based prior-art. More importantly, our new center-surround algorithm is strongly related to the Jacobi iteration that recovers image brightnesses from its Laplacian. Here we showed that Retinex doesn't try and recover the original image but rather attempts to only partially reintegrate the image. The implications of viewing Retinex as a theory of partial reintegration will be investigated in future studies.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council and AgriFoRwArdS CDT, grants EP/S023917/1. and EP/S028730.

References

- [1] C. Bloch. *The HDR1 handbook*. Rockynook, 2007.
- [2] M. Brady and B. K. Horn. Rotationally symmetric operators for surface interpolation. *Computer Vision, Graphics, and Image Processing*, 22(1):70–94, 1983.
- [3] W. L. Briggs, V. E. Henson, and S. F. McCormick. *A multi-grid tutorial*. SIAM, 2000.
- [4] J. Brownlee. A gentle introduction to the rectified linear unit (relu). *Machine learning mastery*, 6, 2019.
- [5] M. D. Fairchild. The hdr photographic survey. In *Color and imaging conference*, pages 233–238. IST, 2007.
- [6] A. Flachot, A. Akbarinia, H. H. Schütt, R. W. Fleming, F. A. Wichmann, and K. R. Gegenfurtner. Deep neural models for color classification and color constancy. *Journal of Vision*, 22(4):17, 1–24, 2022.
- [7] B. V. Funt, F. Ciurea, and J. J. McCann. Tuning retinex parameters. In *Human Vision and Electronic Imaging VII*, pages 358–366. SPIE, 2002.
- [8] C. D. Gilbert and T. N. Wiesel. Functional organization of the visual cortex. *Progress in Brain Research*, 58:209–218, 1983.
- [9] B. K. Horn. Determining lightness from an image. *Computer graphics and image processing*, 3(4):277–299, 1974.
- [10] K. Ji, W. Lei, and W. Zhang. A deep retinex network for underwater low-light image enhancement. *Machine Vision and Applications*, 34(6):122, 2023.
- [11] D. J. Jobson, Z. Rahman, and G. A. Woodell. Properties and performance of a center/surround retinex. *IEEE transactions on image processing*, 6(3):451–462, 1997.
- [12] W. Johnston and E. Kao. Assessment of appearance match by visual observation and clinical colorimetry. *Journal of dental research*, 68(5):819–822, 1989.
- [13] E. H. Land. The retinex. In *Ciba foundation symposium-colour vision: physiology and experimental psychology*, pages 217–227. Wiley Online Library, 1965.
- [14] E. H. Land. An alternative technique for the computation of the designator in the retinex theory of colour vision. *Proceedings of the national academy of sciences*, 83(10):3078–3080, 1986.
- [15] E. H. Land and J. J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.
- [16] M. Lecca. Generalized equation for real-world image enhancement by milano retinex family. *JOSA A*, 37(5):849–858, 2020.
- [17] C. Li, S. Tang, H. K. Kwan, J. Yan, and T. Zhou. Color correction based on cfa and enhancement based on retinex with dense pixels for underwater images. *IEEE Access*, 8: 155732–155741, 2020.
- [18] J. L. Lisani, J. M. Morel, A. B. Petro, and C. Sbert. Analyzing centre/surround retinex. *Information Sciences*, 512: 741–759, 2020.
- [19] M. R. Luo, G. Cui, and B. Rigg. The development of the cie 2000 colour-difference formula: Ciede2000. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur*, 26(5):340–350, 2001.
- [20] D. Marr. *Vision: A computational investigation into the human representation and processing of visual information*. MIT press, 2010.
- [21] J. H. Mathews and K. D. Fink. *Numerical methods using MATLAB*, volume 4. Pearson prentice hall Upper Saddle River, NJ, 2004.
- [22] J. J. McCann. Lessons learned from mondrians applied to real images and colour gamuts. In *Color and imaging conference*, pages 1–8. IST, 1999.
- [23] Z. u. Rahman, D. J. Jobson, and G. A. Woodell. Retinex processing for automatic image enhancement. *Journal of Electronic imaging*, 13(1):100–110, 2004.
- [24] B. C. Regan, C. Julliot, B. Simmen, F. Viénot, P. Charles-Dominique, and J. D. Mollon. Fruits, foliage and the evolution of primate colour vision. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 356(1407):229–283, 2001.
- [25] A. Rizzi and C. Bonanomi. The human visual system described through visual illusions. In *Colour design*, pages 23–41. Elsevier, 2017.
- [26] I. Ruyter, K. Nilner, and B. Möller. Color stability of dental composite resin materials for crown and bridge veneers. *Dental Materials*, 3(5):246–251, 1987.
- [27] M. Stokes. ” a standard default colour space for the internet-srgb”, version 1.10. <http://www.w3.org/pub/WWW/Graphics/Color/sRGB.html>, 1996.
- [28] J. Vazquez-Corral, G. D. Finlayson, and L. Herranz. Improving the perception of low-light enhanced images. *Optics Express*, 32(4):5174–5190, 2024.
- [29] L. Wang, T. Horiuchi, and H. Kotera. High dynamic range image compression by fast integrated surround retinex model. *Journal of Imaging Science and Technology*, 51(1): 34–43, 2007.
- [30] S. S. Yasiran, A. K. Jumaat, A. A. Malek, F. H. Hashim, N. D. Nasriri, S. N. A. S. Hassan, N. Ahmad, and R. Mahmud. Microcalcifications segmentation using three edge detection techniques. In *IEEE (ICEDSA)*, pages 207–211. IEEE, 2012.

Author Biography

Afsaneh Karami is a PhD student at East Anglia University under the supervision of Professor Graham Finlayson. A member of the Colour & Imaging Lab. Also, a member of AgriFoRwArdS CDT.

Professor Graham Finlayson is the Director of the Colour & Imaging Lab. He has published over 200 conference papers, over 75 journal papers and is the inventor of 30+ patents. His interests span perception, colour image processing and physics-based computer vision.