

Individualised Halo-Free Gradient-Domain Colour Image Daltonisation

Ivar Farup ¹ 

¹ Department of Computer Science, Norwegian University of Science and Technology (NTNU), Norway; ivar.farup@ntnu.no

Abstract: Daltonisation refers to the recolouring of images such that details normally lost by colour vision deficient observers become visible. This comes at the cost of introducing artificial colours. In a previous work, we presented a gradient-domain colour image daltonisation method that outperformed previously known methods both in behavioural and psychometric experiments. In the present paper, we improve the method by (i) finding a good first estimate of the daltonised image, thus reducing the computational time significantly, and (ii) introducing local linear anisotropic diffusion, thus effectively removing the halo artefacts. The method uses a colour vision deficiency simulation algorithm as an ingredient, and can thus be applied for any colour vision deficiency, and can even be individualised if the exact individual colour vision is known.

Keywords: daltonisation, colour vision deficiencies, anisotropic diffusion

1. Introduction

There are three types of cone cells, or cones, each with different pigment, in the retinas of the human eye [1]. The cones are referred to as L-cones, M-cones, and S-cones, suggesting sensitivities to long, medium, and short wavelengths, respectively. Among colour normal observers (trichromats) there are only minor variations in the spectral sensitivities of the three cone types. However, about eight percent of the male population and 0.4 percent of the female population suffer from a congenital colour vision deficiency (CVD) [2]. This is manifested either as a change in one of the cone types (anomalous trichromacy), or as a complete lack of one (dichromacy) or, very seldomly, two (monochromacy) of the cone types. Referring to the affected L, M, and S-cones with the Greek prefixes prot-, deuter-, and tri-, respectively, these CVDs can be further characterised as protanomaly, deuteranomaly, and tritanomaly for the anomalous trichromacies, and protanopia, deuteranopia, and tritanopia for the dichromacies. Anomalies in or absence of the L or M cones both lead to various forms of red–green CVDs, whereas anomalies in or absence of the S cones lead to blue–yellow CVDs. In the population, the various forms of red–green CVDs are by far the most common ones.

In many situations, e.g., in graphic design, it can be important for colour normal observers to understand how the world looks through the eyes of a CVD observer. Various simulation methods have been developed for this purpose. In a thorough behavioural comparison of existing CVD simulation methods performed by Simon-Liedtke and Farup [3], the method of Brettel et al. [4] was found to be the best performing one for dichromats.

For the most part, people with CVDs can live completely normal lives. However, in some situations, they may experience certain disadvantages [5]. A considerable effort has been undertaken for coming up with techniques to improve the ability of CVD people to differentiate colour better through optical filters,¹ introduction of lost information through sensory substitution devices [6], and even a medical ‘cure’ has been proposed [7]. For the display of digital images, so-called daltonisation techniques, i.e., recolouring of images for increased visibility for CVD observers, have been proposed [8–

¹ <https://enchroma.com/>, visited 26th September 2020.

11]. CVD simulation algorithms are often an important ingredient in the daltonisation algorithms. Most daltonisation methods use some form of a global recolouring scheme, but daltonisation methods highlighting only colour edges and details that are lost to CVD observers have also been proposed [12]. In a behavioural experiment, the methods of Anagnostopoulos et al. [8] and Kotera [10] were found to be the best performing ones available at the time [13]. A thorough review of global daltonisation methods, i.e., mappings in colour space not taking the local image content into account, was recently given by Ribeiro and Gomes [14]. For image-dependent daltonisation, recent contributions include Chatzistamatis et al.'s method for recoloring of art paintings guided by semantic segmentation [15], and Iqbal et al.'s use of sharpened wavelet bands and inpainting [16].

Recently, we developed a daltonisation method that performed better or on par with state-of-the-art algorithms in both behavioural and psychometric experiments [17]. (An at least principally very similar method, although significantly differently formulated, was almost simultaneously proposed by Zhu et al. [18].) The method was based on a global transformation in the gradient domain of the original colour image, followed by a multi-scale reintegration to obtain the final daltonised image. Although giving good results, the method had two main problems: (i) it was very computationally expensive, and (ii) it tended to produce visual halo artefacts near strong chromatic edges. In the present paper, we present a solution to these two problems by (i) finding a much better initial condition for the iterative algorithm, and (ii) introducing local linear anisotropic diffusion in place of the non-local isotropic diffusion used in the previous method. Together this leads to significantly shorter computing times, as well as to the elimination of visual halo artefacts while still preserving the qualitative behaviour of the previously proposed algorithm.

2. Background

In the previously proposed algorithm [17], the original image was represented as the vector-valued function

$$\mathbf{u}_0 : \Omega \rightarrow \mathcal{C},$$

where $\Omega \subset \mathbb{R}^2$ was the spatial image domain, and $\mathcal{C} \subset \mathbb{R}^3$ the colour space. The spatial gradient $\nabla \mathbf{u}_0$ represented the two spatial directions for all image channels [19]. The CVD simulated image using the method of Brettel et al. [4], was denoted $s(\mathbf{u}_0)$. The difference image between the original image and its simulation, $\mathbf{d}_0 = \mathbf{u}_0 - s(\mathbf{u}_0)$, contained the information that is lost for CVD observers.

We applied a principal component analysis (PCA) [20] of \mathbf{d}_0 to obtain its first principal component $p_1(\mathbf{d}_0)$, representing the main direction of difference, i.e., the main direction of information loss for CVD observers. An orthonormal basis for the daltonisation was built using the Gram–Schmidt process starting from the normalised grey axis taken as either the diagonal of the RGB space, $\mathbf{e}_l = [1, 1, 1]$, or as the direction orthogonal to the plane of constant lightness in, e.g., sRGB space, $\mathbf{e}_l = [0.2126, 0.7152, 0.0722]$ [21], and $p_1(\mathbf{d}_0)$. This resulted in the orthonormal triplet \mathbf{e}_l , \mathbf{e}_d , and \mathbf{e}_c , representing the direction of lightness, difference, and maximally visible chroma change of the CVD observer, respectively.

The gradient of the difference between the original image and the result of the simulation algorithm was then approximated by the projection of the image gradient onto \mathbf{e}_d , $\nabla \mathbf{d}_0 \approx (\nabla \mathbf{u}_0 \cdot \mathbf{e}_d) \mathbf{e}_d$ (for a simulation algorithm that is an orthogonal projection in an equi-luminant plane of the colour space, the approximation is exact). This difference was rotated and scaled into the direction of maximal chromatic visibility, \mathbf{e}_c , and added to the original gradient field to obtain a new gradient tensor

$$\mathbf{G} = \nabla \mathbf{u}_0 + \chi (\nabla \mathbf{u}_0 \cdot \mathbf{e}_d) \mathbf{e}_c, \quad (1)$$

where χ was a suitable scalar field.

In order to determine χ , we assumed the norm of the simulated tensor $S(\mathbf{G})$ to be equal to the norm of the original gradient because the visible gradient for a colour-deficient observer looking at the

daltonised image should equal in magnitude the gradient for a normal-sighted observer looking at the original image,

$$\|S(\mathbf{G})\|_F^2 = \|\nabla \mathbf{u}_0\|_F^2. \quad (2)$$

Here, $\|\dots\|_F^2$ represents the Frobenius norm over both the colour space and the spatial gradient space simultaneously. For simplicity, we approximated the simulation method by a linear projection onto the $(\mathbf{e}_c, \mathbf{e}_l)$ plane, $S(\mathbf{G}) \approx \nabla s(\mathbf{u}_0) + \chi(\nabla \mathbf{u}_0 \cdot \mathbf{e}_d)\mathbf{e}_c$. Inserting this into Equation (2) gives

$$\begin{aligned} & \chi^2 \|(\nabla \mathbf{u}_0 \cdot \mathbf{e}_d)\mathbf{e}_c\|_F^2 \\ & + 2\chi(\nabla \mathbf{u}_0 \cdot \mathbf{e}_d)\mathbf{e}_c : \nabla s(\mathbf{u}_0) \\ & + \|\nabla s(\mathbf{u}_0)\|_F^2 - \|\nabla \mathbf{u}_0\|_F^2 = 0, \end{aligned}$$

a quadratic equation for χ with, in general, two solutions, χ_{\pm} . Which one to use was found by integrating both solutions over the image, selecting the smaller one in absolute value.

Finally, the daltonised image from the tensor \mathbf{G} was found by reintegrating the gradient field iteratively through gradient descent [22] using the finite difference method with explicit time integration – hence referred to as the *isotropic* diffusion equation –

$$\frac{\partial \mathbf{u}}{\partial t} = \nabla \cdot (\nabla \mathbf{u} - \mathbf{G}), \quad (3)$$

with a suitable stopping criterion. Further fine tuning was made using a fidelity term to preserve memory colours such as skin tones.

The original image \mathbf{u}_0 was used as an initial condition for the gradient descent. Since the daltonisation process in some cases lead to a significant change of the colours of the image, and since high image resolution gave very long diffusion paths, the algorithm took very long to converge. In order to reduce the computing time, a multi-scale pyramidal technique was used, first solving the daltonisation problem for a very small down-scaled image, then iteratively using the result as a starting point for the solution at a higher resolution until the solution for the full scale image was found.

3. Proposed Method

3.1. The Initial Value

One main problem of the previously proposed algorithm is the computation time, mainly stemming from long diffusion paths and the use of the original image \mathbf{u}_0 as the initial value. Like for all optimisation problems, the selection of a good initial guess can be crucial for the time of convergence.

Instead of using the original colour image \mathbf{u}_0 as the initial value, we propose here to perform an operation similar to the one applied above to the gradient directly in the colour domain. Letting $s(\mathbf{u}_0)$ denote the simulated image, and constructing \mathbf{e}_l , \mathbf{e}_d , and \mathbf{e}_c as before, the colour information lost by the CVD observer represented by the simulation algorithm can be approximated by the projection of the original image onto \mathbf{e}_d , i.e., $\mathbf{d}_0 \approx (\mathbf{u}_0 \cdot \mathbf{e}_d)\mathbf{e}_d$. Rotating this estimate of the lost information into the direction of maximal chromatic visibility and clipping to the colour gamut gives us a much better first estimate of the daltonised image,

$$\mathbf{u}_s = g(\mathbf{u}_0 + (\mathbf{u}_0 \cdot \mathbf{e}_d)\mathbf{e}_c), \quad (4)$$

where $g(\cdot)$ is the gamut clipping operator. Actually, this simple operation works well as a *simple daltonisation method* all by itself. It produces a recolouring similar to the one by the previously proposed algorithm, but destroys quite some image features and textures due to the gamut clipping operation. However, no halo artefacts are introduced. See Section 4 for sample results.

3.2. Local Linear Anisotropic Diffusion

The second significant problem of the previously proposed method is the production of visible halo artefacts near strong chromatic edges in the images. Switching from non-local isotropic to local linear anisotropic diffusion has proven to be a strong method to reduce haloing problems in various PDE based colour imaging solutions such as colour gamut mapping [23], colour-to-greyscale image conversion [24], and colour image demosaicing [25]. Thus, we propose to apply the same technique here.

The components of the structure tensor \mathbf{S} of the original image can be expressed as [26]

$$S_{ij} = \sum_{\mu, \nu} \frac{\partial u_0^\mu}{\partial x^i} \frac{\partial u_0^\nu}{\partial x^j}.$$

The eigenvalues of \mathbf{S} are denoted λ^+ and λ^- , and the corresponding normalised eigenvectors \mathbf{e}^+ and \mathbf{e}^- are stored as columns in the orthonormal eigenvector matrix \mathbf{E} , such that the structure tensor can be written $\mathbf{S} = \mathbf{E}^T \text{diag}(\lambda^+, \lambda^-) \mathbf{E}$.

From this, the diffusion tensor can, in agreement with Sapiro and Ringach [27], be defined as

$$\mathbf{D} = \mathbf{E}^T \text{diag}(d(\lambda^+), d(\lambda^-)) \mathbf{E},$$

where $d(\lambda)$ is a nonlinear diffusion coefficient function,

$$d(\lambda) = \frac{1}{1 + \kappa \lambda^2},$$

and κ is a suitably chosen numeric constant.

Inserting this into Equation (3), gives the *anisotropic* diffusion equation

$$\frac{\partial \mathbf{u}}{\partial t} = \nabla \cdot [\mathbf{D}(\nabla \mathbf{u} - \mathbf{G})] \quad (5)$$

subject to the constraint $\mathbf{u} \in \mathcal{C}$ with the initial condition $\mathbf{u}(t = 0) = \mathbf{u}_s$.

3.3. Further Simplifications

One complicating factor in the previously proposed method is the computation of the scalar field χ and the selection of the best solution for it. Through experimentation, we have found that we obtain at least visually equally good daltonised images by simply setting $\chi = 1$ throughout, i.e., simply removing χ from Equation (1) giving

$$\mathbf{G} = \nabla \mathbf{u}_0 + (\nabla \mathbf{u}_0 \cdot \mathbf{e}_d) \mathbf{e}_c.$$

A second complicating factor in the previously proposed method was the introduction of the data attachment term for certain memory colours. Since the main purpose of daltonisation methods as such is behavioural, i.e., making observers able to perform certain tasks with the images, we propose here to skip that term altogether.

3.4. Implementation

The methods are discretised using the finite difference method with explicit time integration, and implemented in Python using the NumPy, Scipy, and Matplotlib libraries. The code is freely available on GitHub.²

² <https://github.com/ifarup/anisotropic-daltonisation/>



Figure 1. The image dataset used for testing. Highly saturated images from the web with dominant red–green and blue–yellow contrast. All images are available under the CC0 licence.

4. Results

In order to test the proposed algorithm, we gathered a small set of colourful images from the web with dominating red–green and blue–yellow contrasts. All the gathered images are freely available under the CC0 license.³ The images are shown in Figure 1.

Since the algorithm proposed above works for any kind of CVD simulation algorithm, we could have used the dichromat simulation method by Brettel et al. [4] like in [17]. However, in order to test the behaviour for various kinds of dichromats and anomalous trichromats, we develop a simple parametrised version of their method. The simulation is performed directly in the RGB colour space using the two dichromatic matrices,

$$M_{\text{RG}} = \begin{pmatrix} 1/2 & 1/2 & 0 \\ 1/2 & 1/2 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad M_{\text{BY}} = \begin{pmatrix} 1/2 & 0 & 1/2 \\ 0 & 1/2 & 1/2 \\ 1/4 & 1/4 & 1/2 \end{pmatrix},$$

for red–green (RG) and blue–yellow (BY) dichromats, respectively. In order to simulate also anomalous trichromats, we weigh in the original image with a certain weight α

$$s(\mathbf{u}) = \alpha M\mathbf{u} + (1 - \alpha)\mathbf{u},$$

where M is either M_{RG} or M_{BY} , such that $\alpha = 0$ corresponds to colour normal observers, $0 < \alpha < 1$ to anomalous trichromats, and $\alpha = 1$ to dichromats.

We run the images through the three following daltonisation algorithms:

1. the simple daltonisation algorithm proposed in Equation (4),
2. the isotropic daltonisation algorithm proposed in Equation (3), which is essentially the same as the one proposed in [17], but with the simplifications described in Section 3.3, and the simple daltonisation as the initial condition, and
3. the anisotropic daltonisation algorithm proposed in Equation (5).

This is performed for four different simulated CVD observers: RG and BY with $\alpha = 1.0$ and $\alpha = 0.8$, resulting in a total of $14 \times 3 \times 4 = 504$ daltonised images. Typical results for one original image and one observer (RG dichromat with $\alpha = 1$) are shown in Figure 2.

Zooming into a detail of the sample image, we see in Figure 3 that the simple daltonisation algorithm leads to a significant loss of details. This is caused by the clipping of colours that end up being mapped outside the RGB gamut. The anisotropic daltonisation, on the other hand, is able to preserve much more of the image details. The difference can be quantified by looking at the gamut

³ <https://creativecommons.org/share-your-work/public-domain/cc0/>




Figure 2. Sample image for an RG $\alpha = 1$ dichromatic observer. Top to bottom, left hand column: original, simple daltonisation, isotropic daltonisation, and anisotropic daltonisation. Corresponding simulated images for a the RG $\alpha = 1$ dichromatic observer in the right hand column.



Figure 3. Loss of details in the simple daltonisation algorithm for an RG $\alpha = 1$ dichromatic observer. Left to right: original image, simple daltonisation, and anisotropic daltonisation.

pixel fraction (*GPF*), i.e., the fraction of the image pixels that reside on the gamut boundary, indicated by having pixel values of 0 or 1,

$$GPF = \frac{\#(\text{pixels on the gamut boundary})}{\#(\text{pixels in the image})}.$$

The resulting *GPF* values for all the images are shown as a box plot in Figure 4. Although the distributions for isotropic and anisotropic daltonisation are overlapping, they are statistically significantly different according to the Wilcoxon signed-rank test with a p-value $p = 6.7 \times 10^{-8}$, and with a median in favour of the anisotropic daltonisation.

As noted already in the previous paper [17], visible halo artefacts are observed for the isotropic daltonisation method in many of the images. One such example is shown in Figure 5. The anisotropic daltonisation, on the other hand, produces far less halo artefacts. This can be quantified by comparing how similar the various daltonised images are to the original images by measuring the *PSNR* values,

$$PSNR = -10 \log(MSE),$$

$$MSE = \frac{1}{MN} \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} [I(i, j) - K(i, j)]^2,$$

where I and K indicate the two images that are compared, and M and N denote the horizontal and vertical image resolutions, respectively. The resulting *PSNR* values for all the images are shown as a box plot in Figure 6. Although the distributions for isotropic and anisotropic daltonisation are overlapping, they are statistically significantly different according to the Wilcoxon signed-rank test with a p-value of $p = 7.5 \times 10^{-11}$, and with a median in favour of the anisotropic daltonisation.

5. Conclusion

We have developed a new daltonisation algorithm based on the previously developed method presented in [17]. Unlike the previous method, the proposed method is individualised in the sense

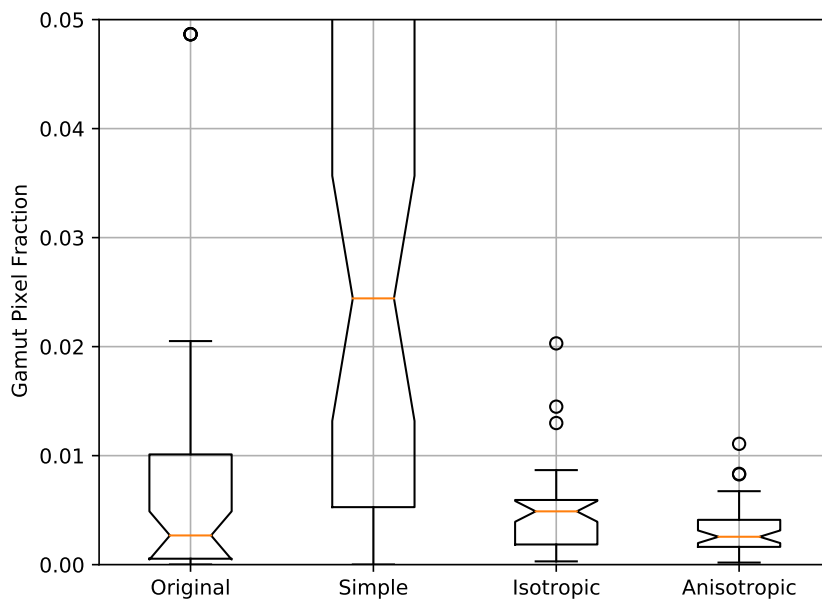


Figure 4. Box plot of the gamut pixel fraction (*GPF*) for the original and the simple, isotropic, and anisotropic daltonised images.



Figure 5. Halo artefacts in the isotropic daltonisation algorithm for a BY $\alpha = 1$ observer. Left to right: original image, isotropic daltonisation, and anisotropic daltonisation.




Figure 6. PSNR (dB) of the daltonised versus the original images for the simple, isotropic, and anisotropic daltonisation algorithms.

that it can be used successfully with any CVD simulation method for dichromats and anomalous trichromats. The method is computationally much more efficient than the previously proposed method due to the simple daltonisation algorithm that acts as a much improved initial value for the iterative numerical scheme. The resulting images are free of the halo artefacts seen previously due to the introduction of a local linear anisotropic diffusion replacing the non-local isotropic diffusion of the previous method. As such, it is the first spatially dependent artefact-free daltonisation algorithm known to the authors. Comparison of the proposed algorithm with other recent state-of-the-art daltonisation methods in behavioural and psychometric experiments is left for future research.

Funding: This research was funded by the Research Council of Norway over the project ‘Individualised Colour Vision-based Image Optimisation’, grant number 287209.

Conflicts of Interest: The author declares no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Valberg, A. *Light Vision Color*; John Wiley & Sons, 2005.
2. Hansen, E. *Fargeblindhet: årsaker, forekomst og typer*; Gyldendal akademisk, 2010.
3. Simon-Liedtke, J.T.; Farup, I. Using a Behavioral Match-to-Sample Method to Evaluate Color Vision Deficiency Simulation Methods. *Journal of Imaging Science and Technology* **2016**, *60*, 50409–1.
4. Brettel, H.; Viénot, F.; Mollon, J.D. Computerized simulation of color appearance for dichromats. *JOSA A* **1997**, *14*, 2647–2655.
5. Simon-Liedtke, J.; Farup, I. Empirical Disadvantages for Color-Deficient People. Proceedings of the AIC2015 Tokyo Color and Image Conference. International Colour Association (AIC), 2015, pp. 391–394.
6. Ward, J.; Wright, T. Sensory substitution as an artificially acquired synaesthesia. *Neuroscience & Biobehavioral Reviews* **2014**, *41*, 26–35.
7. Mancuso, K.; Hauswirth, W.W.; Li, Q.; Connor, T.B.; Kuchenbecker, J.A.; Mauck, M.C.; Neitz, J.; Neitz, M. Gene therapy for red–green colour blindness in adult primates. *Nature* **2009**, *461*, 784–787.

8. Anagnostopoulos, C.N.; Tsekouras, G.; Anagnostopoulos, I.; Kalloniatis, C. Intelligent modification for the daltonization process of digitized paintings. *International Conference on Computer Vision Systems: Proceedings (2007)*, 2007.
9. Huang, J.B.; Chen, C.S.; Jen, T.C.; Wang, S.J. Image recolorization for the colorblind. *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2009, pp. 1161–1164.
10. Kotera, H. Optimal daltonization by spectral shift for dichromatic vision. *Color and Imaging Conference. Society for Imaging Science and Technology*, 2012, pp. 302–308.
11. Kuhn, G.R.; Oliveira, M.M.; Fernandes, L.A. An efficient naturalness-preserving image-recoloring method for dichromats. *IEEE transactions on visualization and computer graphics* **2008**, *14*, 1747–1754.
12. Simon-Liedtke, J.; Farup, I.; Eschbach, R. On the edge: A scalable daltonization method focusing chromatic edges and contrast. *Electronic Imaging* **2017**, *2017*, 28–35.
13. Simon-Liedtke, J.T.; Farup, I. Evaluating color vision deficiency daltonization methods using a behavioral visual-search method. *Journal of Visual Communication and Image Representation* **2016**, *35*, 236–247.
14. Ribeiro, M.; Gomes, A.J. Recoloring algorithms for colorblind people: A survey. *ACM Computing Surveys (CSUR)* **2019**, *52*, 1–37.
15. Chatzistamatis, S.; Rigos, A.; Tsekouras, G.E. Image Recoloring of Art Paintings for the Color Blind Guided by Semantic Segmentation. *International Conference on Engineering Applications of Neural Networks*. Springer, 2020, pp. 261–273.
16. Iqbal, M.; Riaz, M.M.; Ghafoor, A.; Ahmad, A. Image recoloring using specular decomposition and in-painting. *The Imaging Science Journal* **2020**, *68*, 188–194.
17. Simon-Liedtke, J.T.; Farup, I. Multiscale Daltonization in the Gradient Domain. *Journal of Perceptual Imaging* **2018**, *1*, 010503–1–010503–12.
18. Zhu, Z.; Toyoura, M.; Go, K.; Fujishiro, I.; Kashiwagi, K.; Mao, X. Naturalness-and information-preserving image recoloring for red–green dichromats. *Signal Processing: Image Communication* **2019**, *76*, 68–80.
19. Gonzalez, R.C.; Woods, R.E. *Digital image processing*, 3rd ed.; Pearson Education Inc.: Upper Saddle River, NJ, USA, 2008.
20. Jolliffe, I. *Principal component analysis*; Wiley Online Library, 2002.
21. International Electrotechnical Commission (IEC). IEC 61966-2-1:1999. Technical report, International Electrotechnical Commission (IEC), 1999.
22. Rudin, L.I.; Osher, S.; Fatemi, E. Nonlinear total variation based noise removal algorithms. *Phys. D* **1992**, *60*, 259–268.
23. Alsam, A.; Farup, I. Spatial Colour Gamut Mapping by Means of Anisotropic Diffusion. *Computational Colour Imaging Workshop (CCIW)*; Springer: Berlin, 2011; Vol. 6626, *Lecture Notes in Computer Science*, pp. 113–124.
24. Farup, I.; Pedersen, M.; Alsam, A. Colour-to-Greyscale Image Conversion by Linear Anisotropic Diffusion of Perceptual Colour Metrics. *Colour and Visual Computing Symposium*. IEEE, 2018.
25. Thomas, J.B.; Farup, I. Demosaicing of Periodic and Random Colour Filter Arrays by Linear Anisotropic Diffusion. *Journal of Imaging Science and Technology* **2018**, *62*, 050401–1–8.
26. Di Zenzo, S. A note on the gradient of a multi-image. *Computer vision, graphics, and image processing* **1986**, *33*, 116–125.
27. Sapiro, G.; Ringach, D.L. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Transactions on image processing* **1996**, *5*, 1582–1586.