Blue Shift Assumption: Improving Illumination Estimation Accuracy for Single Image from Unknown Source

Nikola Banić and Sven Lončarić
Image Processing Group, Department of Electronic Systems and Information Processing,
Faculty of Electrical Engineering and Computing, University of Zagreb, 10000 Zagreb, Croatia

Keywords: Chromaticity, Color Constancy, Blue, Illumination Estimation, White Balancing.

Abstract: Color constancy methods for removing the influence of illumination on object colors are divided into statistics-based and learning-based ones. The latter have low illumination estimation error, but only on images taken with the same sensor and in similar conditions as the ones used during training. For an image taken with an unknown sensor, a statistics-based method will often give higher accuracy than an untrained or specifically trained learning-based method because of its simpler assumptions not bounded to any specific sensor. The accuracy of a statistics-based method also depends on its parameter values, but for an image from an unknown source these values can be tuned only blindly. In this paper the blue shift assumption is proposed, which acts as a heuristic for choosing the optimal parameter values in such cases. It is based on real-world illumination statistics coupled with the results of a subjective user study and its application outperforms blind tuning in terms of accuracy. The source code is available at http://www.fer.unizg.hr/ipg/resources/color_constancy/.

1 INTRODUCTION

Color constancy enables the human visual system to recognize object colors even under various illumination (Ebner, 2007). Digital cameras also implement some form of computational color constancy (Kim et al., 2012). It first estimates the scene illumination and then it corrects the colors through chromatic adaptation (Gijsenij et al., 2011). The image formation model commonly used for illumination estimation and written under Lambertian assumption is (Gijsenij et al., 2011)

$$f_c(\mathbf{x}) = \int_{\omega} I(\lambda,\mathbf{x})R(\lambda,\mathbf{x})\rho_c(\lambda)d\lambda \qquad (1)$$

where $c \in \{R,G,B\}$ is a color channel of the image $\mathbf{f}$, $\mathbf{x}$ is a given image pixel, λ is the wavelength of the light, ω is the visible spectrum, $I(\lambda,\mathbf{x})$ is the spectral distribution of the light source, $R(\lambda,\mathbf{x})$ is the surface reflectance, and $\rho_c(\lambda)$ is the camera sensitivity of color channel c. Removing $\mathbf{x}$ from $I(\lambda,\mathbf{x})$ by assuming uniform illumination simplifies the problem so that observed light source is then

$$\mathbf{e} = \begin{pmatrix} e_R \\ e_G \\ e_B \end{pmatrix} = \int_{\omega} I(\lambda)\rho(\lambda)d\lambda. \qquad (2)$$

A successful chromatic adaptation requires only the direction of $\mathbf{e}$ (Barnard et al., 2002). However, since both $I(\lambda)$ and $\rho(\lambda)$ are unknown and only $\mathbf{f}$ is given, calculating $\mathbf{e}$ is an ill-posed problem, which is solved by introducing various assumptions. Over time this gave rise to two main groups of illumiantion estimation methods. In the first group are low-level statistics-based methods such as White-patch (Land, 1977; Funt and Shi, 2010) and its improvements (Banić and Lončarić, 2013; Banić and Lončarić, 2014a; Banić and Lončarić, 2014b), Gray-world (Buchsbaum, 1980), Shades-of-Gray (Finlayson and Trezzi, 2004), Grey-Edge (1st and 2nd order) (Van De Weijer et al., 2007). The second group consists of learning-based methods such as gamut mapping (Finlayson et al., 2006), natural image statistics (Gijsenij and Gevers, 2007), spatio-spectral learning (Chakrabarti et al., 2012), simplifying the illumination solution space in various ways (Banić and Lončarić, 2015a; Banić and Lončarić, 2015b; Banić and Lončarić, 2015b; Banić and Lončarić, 2017), using color/edge moments (Finlayson, 2013), regression trees with simple features from color distribution statistics (Cheng et al., 2015), spatial localization (Barron, 2015; Barron and Tsai, 2017), using various convolutional neural network architectures (Bianco et al., 2015; Shi et al., 2016; Hu et al., 2017; Qiu et al., 2018).

Banić, N. and Lončarić, S.
Blue Shift Assumption: Improving Illumination Estimation Accuracy for Single Image from Unknown Source.
DOI: 10.5220/0007394101910197
In *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2019)*, pages 191-197
ISBN: 978-989-758-354-4

The most recent learning-based methods outperform the statistics-based ones by far and in some cases the estimation error can only be attributed to violation of uniform illumination assumption or wrong ground-truth illumination (Zakizadeh et al., 2015). Nevertheless, learning-based methods work so accurately only for images taken with the same sensor and in similar conditions as the ones used in the training dataset. For images taken with another sensor they will usually fail because different sensor characteristics described by $\rho(\lambda)$ were present during image formation (Banić and Lončarić, 2017). Thus, for a single image in the wild from an unknown source learning-based methods will usually not be useful. In cases like this statistics-based methods will often be more accurate because of their simpler assumptions not bounded to any specific sensor. The accuracy of statistics-based methods also depends on their parameter values, but for a single image from an unknown source these values can be tuned only blindly. In this paper the blue shift assumption is proposed, which acts as a heuristic for choosing the optimal parameter values in such cases. It is based on real-world illumination statistics coupled with the results of a subjective user study and it is more accurate then blind tuning.

The paper is structured as follows: Section 2 gives the motivation for making a new assumption, in Section 3 the so called blue shift assumption is proposed, Section 4 contains the experimental results, and Section 5 concludes the paper.

2 MOTIVATION

2.1 Real-world Observations

As mentioned in the introduction, assumptions are needed to handle the illumination estimation problem. While learning-based methods try to extract additional information about images to obtain more accurate illumination estimations, such learning is not possible when only a single image from an unknown source is given and in such cases statistics based methods should be preferred. Most of these methods also have parameters whose values affect the accuracy, but without having any other images from the same source, it is hard to automatically tell which parameter values will give the most accurate result. One of the solutions is to look for additional properties that are usually encountered in natural images. If the basic statistics of illuminations that influence natural images are observed, some general patterns can be observed. For example, by looking at the red chromaticities of systematically measured real-world illumination colors as shown in Fig. 1, it can be seen that the majority of them is centered around lower values. From theoretical aspect this could mean that the real-world illumination or at least the illumination in images of scenes mostly taken by people tends to have lower red chromaticity values i.e. higher blue chromaticity values because of the strong linear connection between the two (Banić and Lončarić, 2015a).

The root cause for this asymmetry can be found in the fact that most of images are taken in outdoor conditions. For example, in the GreyBall dataset (Ciurea and Funt, 2003) that contains 11346 images roughly 57% of them were taken outdoor, for the eight NUS datasets (Cheng et al., 2014) this amounts to roughly 74%, while for RAISE, the challenging real-world image dataset (Dang-Nguyen et al., 2015), this goes over 85%. Namely, in outdoor conditions the most common two illumination sources are the sun and the light scattered across the sky with the latter one having more blue chromaticity and less red chromaticity almost by definition. In practice this means that the scene illumination of taken images will generally tend to be shifted more to the blue.

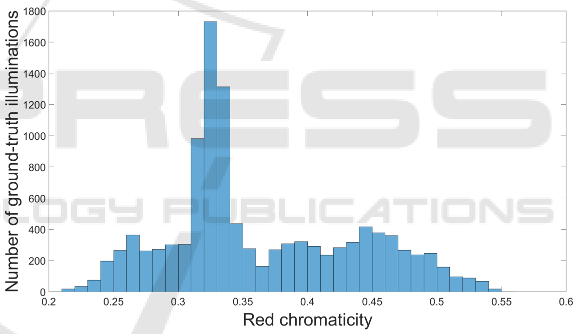

Figure 1: The red chromaticity distribution of the ground-truth illuminations from the GreyBall dataset (Ciurea and Funt, 2003).

2.2 Numerical Observations

Since illumination estimations of statistics-based methods appear "to correlate *roughly* with the actual illuminant" (Finlayson, 2013) i.e. "they occupy roughly the same region in the chromaticity plane" (Banić and Lončarić, 2017), this further means that this empirical information could be applied to illumination estimations. For example, when a statistics-based method produces different results for various parameter values and it has to be decided which one to select as the final one without having any other information, the ones shifted more to the blue should on average probably be preferred. To check to what degree illumination estimations correlate to the actual ground-truth illumination with respect to the being shifted more or less to the blue, it is enough to perform simple counting of illumination estimations given by chosen methods where the red chromaticity is less than the red chromaticity of the ground-truth. Probably the best known statistics-based methods with at least one parameter are Shades-of-Gray, General Gray-World, 1st-order Gray-Edge, and 2nd-order Gray-Edge. All of these methods have the Minkowski norm p parameter, while all but Shades-of-Gray additionally have the σ parameter for Gaussian smoothing. For the purpose of counting the parameter values were constrained to $p \in \{1,\ldots,10\}$ and $\sigma \in \{1,2,3\}$. If for all combinations of these parameter values every of the mentioned methods is applied to each image in the Cube dataset (Banić and Lončarić, 2017), eight NUS datasets (Cheng et al., 2014), and the GreyBall dataset (Ciurea and Funt, 2003), the percentages of the illumination estimations whose red chromaticity is less than the red chromaticity of the ground-truth illumination for a given image are given in Table 1. It can be seen that across methods and datasets in most cases illumination estimations are on average shifted too much to the red.

3 THE PROPOSED ASSUMPTION

3.1 Statement

The previous section boils down to the following two observations: 1) real-world images are taken under illumination that is on average shifted more to the blue, and 2) most commonly used statistics-based illumination estimation methods give illumination estimations that are on average shifted more to the red.

Based on these two observations, the so called **blue shift assumption** can be proposed: *among several candidate illumination estimations for a given image, the ones with the lower red chromaticity are more accurate*. It has to be stressed again that this is only an assumption like e.g. the Gray-World assumption that assumes the average scene reflectance to be achromatic. Such and similar assumptions are often violated, but as as explained in the introduction, they are still required because of the ill-posed nature of the illumination estimation problem. The blue shift assumption is applied to existing illumination estimations, which means that it can be used only in combination with other assumptions used by the methods that initially created these illumination estimations.

3.2 Application

The simplest application of the proposed blue shift assumption to a set of illumination estimations would be to choose the one with the lowest red chromaticity. However, empirically it has been found out that it is better to take the illumination estimation with the second lowest red chromaticity. This founding can be atrributed to the fact that the lowest red chromaticity has a higher probability of being an outlier and should therefore be avoided. Additionally, when inspecting cases where there was no significant outlying, the difference between the lowest and second lowest red chromaticity was not found to be significantly high. All this justifies taking the second lowest red chromaticity in order to avoid potential outliers. The formal notation of the described procedure is given in Algorithm 1.

Algorithm 1: Blue shift assumption.

Input: estimation chromaticities $\mathbb{E} = \{\mathbf{e}^{(1)},\ldots,\mathbf{e}^{(n)}\}$

Output: assumed optimal illumination estimation $\mathbf{e}^*$

$r = \min_i e_R^{(i)}$ ▷ smallest red chromaticity

$m = \arg\min_i \{e_R^{(i)} | r < e_R^{(i)}\}$ ▷ index of second smallest

$\mathbf{e}^* \leftarrow \mathbf{e}^{(m)}$

3.3 Subjective Error Assessment

When the application of the blue shift assumption mistakenly increases the angular error, the resulting chromatically adapted image will by definition tend to be subjectively warmer than an average image with the same illumination estimation error. This is because the red component will be reduced, which in turn will result in less reduction of the redish illumination influence and thus subjectively warmer images. A recent user study has shown ”that when the illuminations are distinct, there is a preference for the outdoor illumination to be corrected resulting in warmer final result” (Cheng et al., 2016). In other words even if the application of the blue shift assumption increases the error, subjectively it is still more acceptable than a colder result with the same angular error as can be seen in Figure 2. However, as can be predicted by the results shown in Table 1, a more usual effect of the blue shift assumption application will be to choose the illumination estimation that is less shifted to the blue than other illumination estimations produced by a given method for various parameter values.

Table 1: Percentages of the illumination estimations whose red chromaticity is less than the red chromaticity of the ground-truth illumination.

	Cube dataset	NUS datasets	GreyBall dataset
Shades-of-Gray	27%	38.76%	35.38%
General Gray-World	29.99%	44.48%	34.13%
1st-order Gray-Edge	34.76%	38%	50.32%
2nd-order Gray-Edge	32.86%	34.24%	50.72%

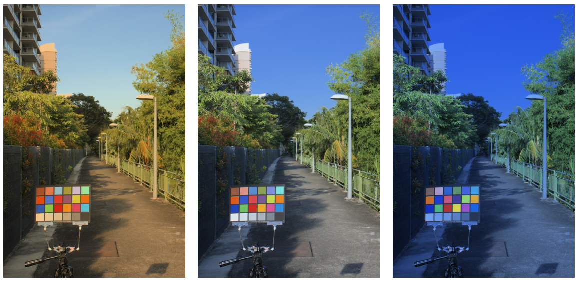

(a) (b) (c)

Figure 2: The effect of red and blue illumination shifting based on the application of the blue shift assumption to the results of the Shades-of-Gray method (Finlayson and Trezzi, 2004) to one of the images from the NUS datasets (Cheng et al., 2014): (a) chromatic adaptation with the result of application of the blue shift assumption and an angular error of 11.56°, (b) chromatic adaptation with ground-truth illumination, and c) chromatic adaptation with the illumination of the same angular error of 11.56° as in a), but with the opposite i.e. red shifting.

4 EXPERIMENTAL RESULTS

4.1 Experimental Setup

The validity of the blue shift assumption was tested on the Cube dataset (Banić and Lončarić, 2017) and eight linear NUS datasets (Cheng et al., 2014) because they all have linear images in accordance with Eq. 1. The ColorChecker dataset was not used because it has been shown on several occasions (Lynch et al., 2013; Finlayson et al., 2017) that it has a public record of biased and wrong usage. Additionally, the GreyBall dataset has also not been used for two reasons. The first one is that it contains non-linear images. The second reasons is that it was used to observe the regularity that serves as the basis for the blue shift assumption, so it may be biased to use it for testing the validity of the assumptions. Among various estimation accuracy measures (Gijsenij et al., 2009; Finlayson and Zakizadeh, 2014; Banić and Lončarić, 2015a) the most commonly used one is the angle between the illumination estimation vector and the ground-truth illumination i.e. the angular error. When describing the angular errors on a dataset by a single statistic, the median angular error is considered to be the best choice (Hordley and Finlayson, 2004) due to the properties of angular error distribution.

4.2 Baseline Methods

The blue shift assumption is supposed to be used in cases when there is only a single image available i.e. when there are no other training images. In such circumstances practically no learning-based method can be either trained or used. As for the statistics-based methods, their parameter values can also not be checked on other images to see which ones should be preferred. A simple baseline method for a single image in this case is to simply average the results obtained for various parameter values without giving preference to any of them. It is also interesting to see what errors are produced by ideally fixed parameter values for a given dataset. While this is definitely unfair because of the advantage of the knowledge of the whole dataset, for comparison purposes it is useful to see how far or close the from such errors is the application of the blue shift assumption.

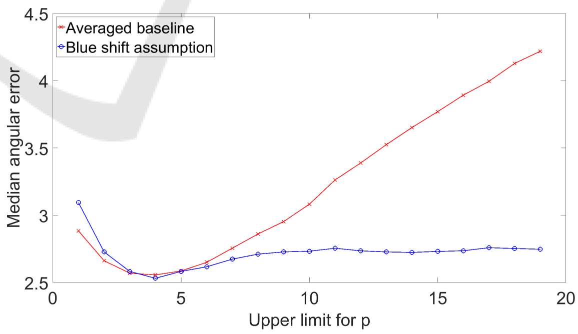

Figure 3: Combining the results of the Shades-of-Gray method on NUS datasets for various upper limits on p.

4.3 Accuracy

The tested statistics-based methods with at least one parameter were as earlier Shades-of-Gray, General Gray-World, 1st-order Gray-Edge, and 2nd-order Gray-Edge. The parameter values were also constrained in the same way. Table 2 shows the angular errors obtained by the baseline methods and the

Table 2: Combined angular errors on the linear images of the Cube dataset (Banić and Lončarić, 2017) and eight NUS datasets (Cheng et al., 2014) (lower is better). The used format is the same as in (Barron and Tsai, 2017).

Algorithm	**Cube dataset**						**NUS datasets**					
	Mean	**Med.**	**Tri.**	**Best 25%**	**Worst 25%**	**Avg.**	**Mean**	**Med.**	**Tri.**	**Best 25%**	**Worst 25%**	**Avg.**
Shades-of-Gray (Finlayson and Trezzi, 2004)												
Averaged baseline	2.65	1.91	2.07	0.42	6.18	1.94	3.65	2.95	3.16	0.95	7.46	3.00
Ideally fixed parameters	2.55	1.72	1.90	0.38	6.14	1.81	3.44	**2.61**	2.78	0.83	7.42	2.74
Blue shift assumption	2.18	**1.51**	1.66	0.36	5.16	1.59	3.82	2.73	2.92	0.91	8.67	2.99
General Gray-World (Barnard et al., 2002)												
Averaged baseline	2.60	1.75	1.93	0.39	6.22	1.84	3.16	**2.35**	2.53	0.70	6.90	2.47
Ideally fixed parameters	2.47	1.56	1.77	0.37	6.15	1.73	3.28	2.43	2.58	0.70	7.34	2.53
Blue shift assumption	2.15	**1.44**	1.59	0.35	5.20	1.55	3.40	2.58	2.71	0.81	7.48	2.70
1st-order Gray-Edge (Van De Weijer et al., 2007)												
Averaged baseline	2.43	1.63	1.83	0.49	5.72	1.82	3.38	2.55	2.74	0.89	7.26	2.74
Ideally fixed parameters	2.40	1.52	1.76	0.45	5.78	1.76	3.07	**2.11**	2.33	0.70	7.05	2.37
Blue shift assumption	2.07	**1.43**	1.59	0.49	4.68	1.61	3.44	2.42	2.60	0.84	7.84	2.70
2nd-order Gray-Edge (Van De Weijer et al., 2007)												
Averaged baseline	2.70	1.93	2.12	0.74	5.97	2.18	3.83	3.00	3.18	1.17	7.90	3.20
Ideally fixed parameters	2.43	**1.53**	1.77	0.46	5.83	1.77	3.11	**2.28**	2.42	0.78	6.91	2.47
Blue shift assumption	2.25	1.68	1.82	0.53	4.90	1.78	3.88	2.63	2.84	0.92	9.03	3.00

blue shift assumption on linear images of the Cube dataset and eight NUS datasets. In all but one case the blue shift assumption leads to higher accuracy than the simple averaged baseline method. As for the ideally fixed parameters, on NUS datasets they always give lower errors and on the Cube dataset only once. An additional property of the blue shift assumption is that it is more stable than the simple averaged baseline method. Namely, as the number of possible parameter values increases, the accuracy of the averaged result tends to decrease, while the one of the blue shift assumption remains much more stable. In Fig. 3 this is shown for the results of applying the Shades-of-Gray method to the NUS datasets. There is another important thing to be observed in Table 2, namely the fact that for supposedly more accurate methods the blue shift assumption also results in higher estimation accuracy. For example in terms of median angular error on both the Cube and the NUS datasets the blue shift assumption gives higher accuracy for 1st-order Gray-Edge than for General Gray-World and it also gives higher accuracy for General Gray-World than for Shades-of-Gray.

When the blue shift assumption is applied to non-linear images, the positive effect is visible to a lesser extent as can be seen on the results for the non-linear images of the GreyBall dataset (Ciurea and Funt, 2003) shown in Table 3. Here the blue shift assumption leads to higher accuracy than the simple averaged baseline method in half of the cases. The main difference between these images and the images from the Cube and NUS datasets is that the images in the GreyBall dataset are non-linear, which practically always leads to higher errors in illumination estimation accuracy for a given method (Gijsenij et al., 2011; Gijsenij et al., 2018). This shows how increased estimation errors also lead to inefficiency of the blue shift assumption. Something similar could also have been observed in Table 2 for linear images where the blue shift assumption was shown to be less efficient when applied to results of less accurate methods.

4.4 Discussion

The experimental results clearly show the benefits of the blue shift assumption over the simple averaged baseline method. In addition to being more accurate, this assumption is also more stable when the number of illumination estimations to be combined changes. The blue shift assumption failed to outperform the averaged baseline only for the General Gray-World method on the NUS datasets. It is also interesting to note that in many cases the blue shift assumption outperforms the results obtained by using ideally fixed parameters, which shows the benefits of the assumption's dynamical parameter values adjustment. The cases for which the blue shift assumption fails are the ones where all underlying illumination estimations are already erroneously shifted to the blue or where they are all very close to the ground-truth illumination. Nevertheless, as explained previously in more detail in Section 3.3, the resulting errors are relatively acceptable.

Table 3: Angular errors on the non-linear images of the GreyBall dataset (Ciurea and Funt, 2003) (lower is better). The used format is the same as in (Barron and Tsai, 2017).

Algorithm	Mean	Med.	Tri.	Best 25%	Worst 25%	Avg.
Shades-of-Gray (Finlayson and Trezzi, 2004)						
Averaged baseline	6.23	5.37	5.58	1.74	12.20	5.25
Ideally fixed parameters	6.11	**5.28**	5.48	1.75	11.88	5.17
Blue shift assumption	6.29	5.40	5.62	1.72	12.35	5.27
General Gray-World (Barnard et al., 2002)						
Averaged baseline	6.50	5.58	5.81	1.78	12.76	5.44
Ideally fixed parameters	6.24	**5.37**	5.60	1.76	12.16	5.26
Blue shift assumption	6.49	5.51	5.79	1.73	12.83	5.40
1st-order Gray-Edge (Van De Weijer et al., 2007)						
Averaged baseline	6.50	5.58	5.81	1.78	12.76	5.44
Ideally fixed parameters	6.24	**5.37**	5.60	1.76	12.16	5.26
Blue shift assumption	6.49	5.51	5.79	1.73	12.83	5.40
2nd-order Gray-Edge (Van De Weijer et al., 2007)						
Averaged baseline	6.82	5.15	5.71	1.43	14.73	5.31
Ideally fixed parameters	6.10	**4.85**	5.28	1.64	12.42	5.02
Blue shift assumption	7.88	5.80	6.49	1.51	17.54	6.01

5 CONCLUSIONS

The so called blue shift assumption has been proposed to increase the accuracy of statistics-based methods for the case when there is only one image given from an unknown sensor. It was experimentally shown to outperform the simple averaging baseline method when no preference is given to any of the parameter values of a chosen statistics-based method. The results of a user study can additionally be used to show that even in failure cases the blue shift assumption produces results that are subjectively more acceptable than average failure cases. In future some better outlier removal strategies in the blue shift assumption will be researched to further increase the accuracy. Another direction will be to look for some other similar properties that can be used when only on a single image is given.

ACKNOWLEDGEMENT

This work has been supported by the Croatian Science Foundation under Project IP-06-2016-2092.

REFERENCES

Banić, N. and Lončarić, S. (2015a). Color Cat: Remembering Colors for Illumination Estimation. *Signal Processing Letters, IEEE*, 22(6):651–655.

Banić, N. and Lončarić, S. (2015b). Using the red chromaticity for illumination estimation. In *Image and Signal Processing and Analysis (ISPA), 2015 9th International Symposium on*, pages 131–136. IEEE.

Banić, N. and Lončarić, S. (2017). Unsupervised Learning for Color Constancy. *arXiv preprint arXiv:1712.00436*.

Banić, N. and Lončarić, S. (2013). Using the Random Sprays Retinex Algorithm for Global Illumination Estimation. In *Proceedings of The Second Croatian Computer Vision Workshop (CCVW 2013)*, pages 3–7. University of Zagreb Faculty of Electrical Engineering and Computing.

Banić, N. and Lončarić, S. (2014a). Color Rabbit: Guiding the Distance of Local Maximums in Illumination Estimation. In *Digital Signal Processing (DSP), 2014 19th International Conference on*, pages 345–350. IEEE.

Banić, N. and Lončarić, S. (2014b). Improving the White patch method by subsampling. In *Image Processing (ICIP), 2014 21st IEEE International Conference on*, pages 605–609. IEEE.

Banić, N. and Lončarić, S. (2015a). A Perceptual Measure of Illumination Estimation Error. In *VISAPP*, pages 136–143.

Banić, N. and Lončarić, S. (2015b). Color Dog: Guiding the Global Illumination Estimation to Better Accuracy. In *VISAPP*, pages 129–135.

Barnard, K., Cardei, V., and Funt, B. (2002). A comparison of computational color constancy algorithms. i: Methodology and experiments with synthesized data. *Image Processing, IEEE Transactions on*, 11(9):972–984.

Barron, J. T. (2015). Convolutional Color Constancy. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 379–387.

Barron, J. T. and Tsai, Y.-T. (2017). Fast Fourier Color Constancy. In *Computer Vision and Pattern Recognition,* *2017. CVPR 2017. IEEE Computer Society Conference on*, volume 1. IEEE.

Bianco, S., Cusano, C., and Schettini, R. (2015). Color Constancy Using CNNs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 81–89.

Buchsbaum, G. (1980). A spatial processor model for object colour perception. *Journal of The Franklin Institute*, 310(1):1–26.

Chakrabarti, A., Hirakawa, K., and Zickler, T. (2012). Color constancy with spatio-spectral statistics. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(8):1509–1519.

Cheng, D., Abdelhamed, A., Price, B., Cohen, S., and Brown, M. S. (2016). Two Illuminant Estimation and User Correction Preference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 469–477.

Cheng, D., Prasad, D. K., and Brown, M. S. (2014). Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *JOSA A*, 31(5):1049–1058.

Cheng, D., Price, B., Cohen, S., and Brown, M. S. (2015). Effective learning-based illuminant estimation using simple features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1000–1008.

Ciurea, F. and Funt, B. (2003). A large image database for color constancy research. In *Color and Imaging Conference*, volume 2003, pages 160–164. Society for Imaging Science and Technology.

Dang-Nguyen, D.-T., Pasquini, C., Conotter, V., and Boato, G. (2015). RAISE: a raw images dataset for digital image forensics. In *Proceedings of the 6th ACM Multimedia Systems Conference*, pages 219–224. ACM.

Ebner, M. (2007). *Color Constancy*. The Wiley-IS&T Series in Imaging Science and Technology. Wiley.

Finlayson, G. D. (2013). Corrected-moment illuminant estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1904–1911.

Finlayson, G. D., Hemrit, G., Gijsenij, A., and Gehler, P. (2017). A Curious Problem with Using the Colour Checker Dataset for Illuminant Estimation. In *Color and Imaging Conference*, volume 2017, pages 64–69. Society for Imaging Science and Technology.

Finlayson, G. D., Hordley, S. D., and Tastl, I. (2006). Gamut constrained illuminant estimation. *International Journal of Computer Vision*, 67(1):93–109.

Finlayson, G. D. and Trezzi, E. (2004). Shades of gray and colour constancy. In *Color and Imaging Conference*, volume 2004, pages 37–41. Society for Imaging Science and Technology.

Finlayson, G. D. and Zakizadeh, R. (2014). Reproduction angular error: An improved performance metric for illuminant estimation. *perception*, 310(1):1–26.

Funt, B. and Shi, L. (2010). The rehabilitation of MaxRGB. In *Color and Imaging Conference*, volume 2010, pages 256–259. Society for Imaging Science and Technology.

Gijsenij, A. and Gevers, T. (2007). Color Constancy using Natural Image Statistics. In *CVPR*, pages 1–8.

Gijsenij, A., Gevers, T., and Lucassen, M. P. (2009). Perceptual analysis of distance measures for color constancy algorithms. *JOSA A*, 26(10):2243–2256.

Gijsenij, A., Gevers, T., and Van De Weijer, J. (2011). Computational color constancy: Survey and experiments. *Image Processing, IEEE Transactions on*, 20(9):2475–2489.

Gijsenij, A., Gevers, T., and van de Weijer, J. (2018). Color Constancy — Research Website on Illuminant Estimation.

Hordley, S. D. and Finlayson, G. D. (2004). Re-evaluating colour constancy algorithms. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, volume 1, pages 76–79. IEEE.

Hu, Y., Wang, B., and Lin, S. (2017). Fully Convolutional Color Constancy with Confidence-weighted Pooling. In *Computer Vision and Pattern Recognition, 2017. CVPR 2017. IEEE Conference on*, pages 4085–4094. IEEE.

Kim, S. J., Lin, H. T., Lu, Z., Süsstrunk, S., Lin, S., and Brown, M. S. (2012). A new in-camera imaging model for color computer vision and its application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(12):2289–2302.

Land, E. H. (1977). *The retinex theory of color vision*. Scientific America.

Lynch, S. E., Drew, M. S., and Finlayson, k. G. D. (2013). Colour Constancy from Both Sides of the Shadow Edge. In *Color and Photometry in Computer Vision Workshop at the International Conference on Computer Vision*. IEEE.

Qiu, J., Xu, H., Ma, Y., and Ye, Z. (2018). PILOT: A Pixel Intensity Driven Illuminant Color Estimation Framework for Color Constancy. *arXiv preprint arXiv:1806.09248*.

Shi, W., Loy, C. C., and Tang, X. (2016). Deep Specialized Network for Illuminant Estimation. In *European Conference on Computer Vision*, pages 371–387. Springer.

Van De Weijer, J., Gevers, T., and Gijsenij, A. (2007). Edge-based color constancy. *Image Processing, IEEE Transactions on*, 16(9):2207–2214.

Zakizadeh, R., Brown, M. S., and Finlayson, G. D. (2015). A Hybrid Strategy For Illuminant Estimation Targeting Hard Images. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 16–23.