

Color Dog: Guiding the Global Illumination Estimation to Better Accuracy

Nikola Banić¹ and Sven Lončarić¹

¹*Image Processing Group, Department of Electronic Systems and Information Processing, Faculty of Electrical Engineering and Computing, University of Zagreb, 10000 Zagreb, Croatia*
{nikola.banic, sven.loncaric}@fer.hr

Keywords: Clustering, color constancy, illumination estimation, image enhancement, white balancing.

Abstract: An important part of image enhancement is color constancy, which aims to make image colors invariant to illumination. In this paper the Color Dog (CD), a new learning-based global color constancy method is proposed. Instead of providing one, it corrects the other methods' illumination estimations by reducing their scattering in the chromaticity space by using a its previously learning partition. The proposed method outperforms all other methods on most high-quality benchmark datasets. The results are presented and discussed.

1 INTRODUCTION

Color constancy is the ability to recognize object colors regardless of the scene illumination (Ebner, 2007). Achieving it is often used as a pre-processing method in image processing because depending on scene illumination, the image colors may differ as shown in Fig. 1. Two steps are needed to achieve computational color constancy: illumination estimation, the essential step, and chromatic adaptation using the estimation, a relatively easy step. Both steps often use the following image f formation model, which includes Lambertian assumption:

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

where c is a color channel, \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, $R(\mathbf{x}, \lambda)$ is the surface reflectance and $\rho_c(\lambda)$ is the camera sensitivity of the c -th color channel. With uniform illumination assumed, \mathbf{x} is removed from $I(\lambda, \mathbf{x})$ and the observed color of the light source \mathbf{e} is:

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

This paper was presented at VISAPP 2015: <http://www.visapp.visigrapp.org/?y=2015>.



Figure 1: The same scene (a) with and (b) without illumination color cast.

For successful chromatic adaptation only the direction of \mathbf{e} is important and its amplitude can be disregarded. As it is often the case that the values of $I(\lambda)$ and $\rho_c(\lambda)$ are unknown, calculating \mathbf{e} is an ill-posed problem and additional assumptions are taken to solve it. This has resulted in many color constancy methods that form at least two groups. The first group is formed of low-level statistics-based methods like White-patch (WP) (Land, 1977) and its improved version (Banić and Lončarić, 2014b), Gray-world (GW) (Buchsbau, 1980), Shades-of-Gray (SoG) (Finlayson and Trezzi, 2004), Grey-Edge (1st and 2nd order (GE1 and GE2)) (Van De Weijer et al., 2007a), Weighted Gray-Edge (Gijssenij et al., 2012), using bright pixels (BP) (Joze et al., 2012), Color Sparrow (CS) (Banić and Lončarić, 2013), Color Rabbit (CR) (Banić and Lončarić, 2014a), using color distribution (CD) (Cheng et al., 2014b). The second group is formed of learning-based methods like gamut mapping (pixel, edge, and intersection based -

PG, EG, and IG) (Finlayson et al., 2006), using neural networks (Cardei et al., 2002), using high-level visual information (HLVI) (Van De Weijer et al., 2007b), natural image statistics (NIS) (Gijssenij and Gevers, 2007), Bayesian learning (BL) (Gehler et al., 2008), spatio-spectral learning (maximum likelihood estimate (SL) and with gen. prior (GP)) (Chakrabarti et al., 2012), exemplar-based learning (EB) (Joze and Drew, 2012), Color Cat (CC) (Banić and Lončarić, 2015). In devices with limited computation power like digital cameras, faster low-level statistics-based methods are used (Deng et al., 2011) because the more accurate learning-based methods are slower.

Recently, a new learning-based method using the illumination statistics has been proposed (Zhang and Batur, 2014), which limits the possible values of \mathbf{e} to only a set of illuminations. The most appropriate illumination for a given image is then selected by means of a classifier that uses the image chromaticity histogram bin values as features. In this paper a new method is proposed, which also selects the most appropriate illumination from a set of illuminations, but the approach is significantly different. The selection is performed by a voting procedure, which uses the illumination estimations of other existing methods. The selected illumination can be interpreted as the correction of used initial illumination estimation. With this correction the very fast, but less accurate statistics-based methods can simply be improved to outperform most of the other state-of-the-art methods in terms of accuracy. At the same time the accuracy of already accurate methods is improved even further.

The paper is structured as follows: In Section II the proposed method is described and in Section III it is tested and the results are presented and discussed.




Figure 2: The rb -chromaticities of the Sony dataset (Cheng et al., 2014b) ground-truth illuminations.

2 THE PROPOSED METHOD

2.1 Motivation

Since it is only the direction of the observed light source \mathbf{e} that matters, chromaticity can be used to describe illumination. The rgb chromaticity components of a color described by its RGB components are calculated by scaling the RGB components so that the equation $r + g + b = 1$ holds:

$$r = \frac{R}{R + G + B}, \quad (3)$$

$$g = \frac{G}{R + G + B}, \quad (4)$$

$$b = \frac{B}{R + G + B}. \quad (5)$$

Each estimation of \mathbf{e} can result in various values. Fig. 2 shows the ground-truth illumination chromaticities for the Sony dataset (Cheng et al., 2014b). It can be seen that there is a certain regularity that can also be seen for any other benchmark dataset. This has been used in (Zhang and Batur, 2014) where the ground-truth illumination chromaticities are clustered by performing the k -means clustering (Vassilvitskii and University, 2007) to obtain the cluster centers, which are used to sparsely represent the possible illumination values. For a given image the cluster center that most appropriately approximates the image scene illumination is chosen by using the image chromaticity histogram and a machine learning algorithm. In this way the problem of illumination estimation is significantly simplified by transforming it into classification problem.




Figure 3: The Canon1 dataset (Cheng et al., 2014b) ground-truth illumination clustering example.

Fig. 3 shows a possible clustering of the Canon1 dataset (Cheng et al., 2014b) ground-truth illuminations. Since the illumination chromaticities are densely placed together, representing any of them with one member of a well chosen small set of illumination chromaticities results in only a small error




Figure 4: The illumination estimation rb -chromaticities on the Canon1 dataset (Cheng et al., 2014b) for (a) White-patch and (b) Gray-world method.

and consequently in a good approximation. This was done in (Zhang and Batur, 2014), but the method described there is a learning-based method that is not very appropriate for devices with limited computational power.

As seen in Fig. 4, the illumination estimations chromaticities of White-patch and Gray-world are scattered around the ground-truth illumination chromaticities. Similar arrangements can be observed for other methods as well. Such arrangements lead to the motivation of trying to correct the methods' illumination estimations by getting their chromaticities closer to the region occupied by the ground-truth illumination chromaticities.

2.2 Realisation

Instead of performing classification by extracting features and applying a machine learning algorithm, we propose a method that chooses the most appropriate center by performing a voting where the voters are some existing illumination estimation methods that cast a vote of different strength for each of the available centers. Since both the methods' illumination estimations and the centers are vectors, the vote that each of the used methods casts for each of the centers can be defined as the cosine of the angle between the center and the method's illumination estimation. The center with the maximum sum of the votes is the proposed method's illumination estimation.

Since dogs are known for their leading abilities and the proposed method leads the voters' illumination estimations to higher accuracy, it was named the Color Dog (CD). The pseudocode for the application phase of the Color Dog method is given in Algorithm 1. An example of Color Dog correction with only one voter is shown in Fig. 5.

The voters are chosen by considering where the illumination estimation is to be applied. On digital cameras it might be better to use statistics-based vot-

Algorithm 1 Color Dog Application

```

1:  $I = GetImage()$ 
2: for all  $voter_i \in \{voter_1, \dots, voter_n\}$  do
3:    $e_i = voter_i.EstimateIllumination(I)$ 
4: end for
5:  $e = \arg \max_{c \in centers} \left( \sum_{i=1}^n \frac{c \cdot e_i}{\|c\| \cdot \|e_i\|} \right)$ 

```

ers since they are fast. If speed is not critical, then correction the illumination estimation of learning-based methods might result in an even higher accuracy. If the Color Dog uses used voters v_1, v_2, \dots, v_n , then this is notated as $CD_{v_1, v_2, \dots, v_n}$.

In addition to any parameters of the voter methods, the center positions also need to be learned. This is done by performing k -means algorithm on the ground-truth illuminations of the learning set. Additionally, what also needs to be determined is the number of centers, which is a hyperparameter and each value represents a different model. More centers result in a more accurate chromaticity space representation and a harder classification problem, so the optimal number of centers has to be chosen carefully. This is done in the model selection process (Japkowicz and Shah, 2011), which conducts a grid search guided by cross-validating the proposed method for a given number of centers. At the end the selected model i.e. number of centers is the one that resulted in the lowest generalization error.

3 EXPERIMENTAL RESULTS

3.1 Benchmark datasets

Since the image formation model used in Eq. (1) is linear and in digital cameras color constancy methods are implemented to work on linear images (Gijsenij




Figure 5: Example of the Color Dog method with one voter choosing the closest center.



Figure 6: Example of chromatic adaptation based on the methods' illumination estimation and respective illumination estimation errors: (a) do nothing, (b) White-patch, (c) Gray-world, (d) Color Distribution, (e) Color Rabbit), and (f) proposed method.

et al., 2011), datasets with linear image were used to test the accuracy of the proposed method in such environments. Until recently the only publicly available and well-known raw-based dataset with linear images was the Shi's and Funt's re-processed linear version (L. Shi, 2014) of the ColorChecker dataset (Gijzenj and Gevers, 2007). However, in most publications this dataset was used without subtracting the black level (Lynch et al., 2013), which led to wrong estimations. Since this might lead to certain comparison problems, the linear ColorChecker was not used. Instead the nine new NUS datasets (Cheng et al., 2014b) were used. Images of these datasets are of high-quality and each dataset corresponds to a different camera. In addition to these datasets, the non-linear GreyBall dataset (Ciurea and Funt, 2003) and its approximated linear version were used because they are the largest available benchmark datasets.

In the scene of each image is a calibration object used to extract the ground-truth illumination of the scene. When an illumination estimation method is applied to the image, the calibration object is first

masked out in order to avoid bias. After the illumination estimation is performed, the angle between the resulting vector and the ground-truth vector is calculated and used as the error measure i.e. as angular error. The commonly used statistics descriptor used to describe a method's performance on a dataset is the median of per image angular error (Hordley and Finlayson, 2004). The mean is less important because the angular error distribution is in many cases non-symmetric.

3.2 Used voters

Because speed is a desirable feature of illumination estimation methods, especially for real-time embedded system implementations, one of the tested Color Dog voter method sets contained two of the simplest methods that have no parameters: White-patch and Gray-world method. They are very fast, but not very accurate and any improvement of their accuracy is significant. White-patch illumination estimation

Table 1: Angular error of selected low-level statistics-based methods, the proposed method, and selected learning-based methods on nine NUS benchmark image databases (lower is better).

	Low-level statistics-based methods							Learning-based methods						
Method	CR	CD	GW	WP	GGW	GE1	GE2	CD _{GW,WP}	PG	EG	IG	ML	GP	NIS
Dataset	Mean angular error (°)													
Canon1	3.09	2.93	5.16	7.99	3.16	3.45	3.47	3.13	6.13	6.07	6.37	3.58	3.21	4.18
Canon2	2.81	2.81	3.89	10.96	3.24	3.22	3.21	2.83	14.51	15.36	14.46	2.80	2.67	3.43
Fuji	2.94	3.15	4.16	10.20	3.42	3.13	3.12	3.36	8.59	7.76	6.80	3.12	2.99	4.05
Nikon1	3.06	2.90	4.38	11.64	3.26	3.37	3.47	3.19	10.14	13.00	9.67	3.22	3.15	4.10
Oly	2.65	2.76	3.44	9.78	3.08	3.02	2.84	2.57	6.52	13.20	6.21	2.92	2.86	3.22
Pan	2.89	2.96	3.82	13.41	3.12	2.99	2.99	2.84	6.00	5.78	5.28	2.93	2.85	3.70
Sam	2.94	2.91	3.90	11.97	3.22	3.09	3.18	2.92	7.74	8.06	6.80	3.11	2.94	3.66
Sony	2.88	2.93	4.59	9.91	3.20	3.35	3.36	2.83	5.27	4.40	5.32	3.24	3.06	3.45
Nikon2	3.57	3.81	4.60	12.75	4.04	3.94	3.95	3.37	11.27	12.17	11.27	3.80	3.59	4.36
Dataset	Median angular error (°)													
Canon1	2.08	2.01	4.15	6.19	2.35	2.48	2.44	1.72	4.30	4.68	4.72	2.80	2.67	3.04
Canon2	1.86	1.89	2.88	12.44	2.28	2.07	2.29	1.85	14.83	15.92	14.72	2.32	2.03	2.46
Fuji	1.84	2.15	3.30	10.59	2.60	1.99	2.00	1.81	8.87	8.02	5.90	2.70	2.45	2.95
Nikon1	1.91	2.08	3.39	11.67	2.31	2.22	2.19	1.94	10.32	12.24	9.24	2.43	2.26	2.40
Oly	1.79	1.87	2.58	9.50	2.15	2.11	2.18	1.46	4.39	8.55	4.11	2.24	2.21	2.17
Pan	1.70	2.02	3.06	18.00	2.23	2.16	2.04	1.69	4.74	4.85	4.23	2.28	2.22	2.28
Sam	1.88	2.03	3.00	12.99	2.57	2.23	2.32	1.89	7.91	6.12	6.37	2.51	2.29	2.77
Sony	2.10	2.33	3.46	7.44	2.56	2.58	2.70	1.77	4.26	3.30	3.81	2.70	2.58	2.88
Nikon2	2.42	2.72	3.44	15.32	2.92	2.99	2.95	2.12	10.99	11.64	11.32	2.99	2.89	3.51
Dataset	Trimean angular error (°)													
Canon1	2.56	2.22	4.46	6.98	2.50	2.74	2.70	2.08	4.81	4.87	5.13	2.97	2.79	3.30
Canon2	2.17	2.12	3.07	11.40	2.41	2.36	2.37	2.07	14.78	15.73	14.80	2.37	2.18	2.72
Fuji	2.13	2.41	3.40	10.25	2.72	2.26	2.27	2.20	8.64	7.70	6.19	2.69	2.55	3.06
Nikon1	2.23	2.19	3.59	11.53	2.49	2.52	2.58	2.14	10.25	11.75	9.35	2.59	2.49	2.77
Oly	2.01	2.05	2.73	9.54	2.35	2.26	2.20	1.72	4.79	10.88	4.63	2.34	2.28	2.42
Pan	2.12	2.31	3.15	14.98	2.45	2.25	2.26	1.87	4.98	5.09	4.49	2.44	2.37	2.67
Sam	2.18	2.22	3.15	12.45	2.66	2.32	2.41	2.05	7.70	6.56	6.40	2.63	2.44	2.94
Sony	2.26	2.42	3.81	8.78	2.68	2.76	2.80	2.03	4.45	3.45	4.13	2.82	2.74	2.95
Nikon2	2.67	3.10	3.69	13.80	3.22	3.21	3.38	2.38	11.11	12.01	11.30	3.11	2.96	3.84

looks like this:

$$\mathbf{e}_{wp} = \begin{pmatrix} \max f_R \\ \max f_G \\ \max f_B \end{pmatrix}, \quad (6)$$

but for better performance, clipped pixels should be excluded from the maxima calculation (Funt and Shi, 2010). Gray-world illumination estimation is performed in the following way:

$$\mathbf{e}_{gw} = \frac{\int \mathbf{f}(x) dx}{\int dx}. \quad (7)$$

The described combination of the voters CD_{WP,GGW} was tested on all used datasets with the goal of examining how much can the accuracy of some of the simplest and least accurate methods be improved with the computational cost almost intact. For the challenging GreyBall dataset the combinations CD_{EB} and CD_{CC} were also tested where the voter methods described in (Joze and Drew, 2012) and (Banić and Lončarić, 2015) were used to examine if the accuracy of already successful methods can be improved even further.

3.3 Accuracy

Table 1 shows the error statistics achieved on the NUS datasets, Table 2 on the original GreyBall dataset, and Table 3 on its linear version. The results for other methods were taken from (Cheng et al., 2014b) (Cheng et al., 2014a) (A. Gijzenj and van de Weijer, 2014). Because the proposed method is a learning one, like for other learning-based methods the 3-fold cross-validation was used on the NUS datasets. On the GreyBall dataset the 15-fold cross-validation was used with the folds provided by the authors. In the loop performing the cross-validation there was another cross-validation to determine the optimal number of centers so that the whole testing was performed using a nested cross-validation.

For the most of the NUS datasets the proposed CD_{WP,GGW} outperforms all other methods in terms of median angular error, which is a significant result. Since the median angular error is for all NUS datasets

Table 2: Performance of different color constancy methods on the original GreyBall dataset (lower is better).

method	mean (°)	median (°)	trimean (°)
do nothing	8.28	6.70	7.25
Low-level statistics-based methods			
GW	7.87	6.97	7.14
WP	6.80	5.30	5.77
SoG	6.14	5.33	5.51
general GW	6.14	5.33	5.51
GE1	5.88	4.65	5.11
GE2	6.10	4.85	5.28
Learning-based methods			
PG	7.07	5.81	6.12
EG	6.81	5.81	6.03
IG	6.93	5.80	6.05
NIS	5.19	3.93	4.31
EB	4.38	3.43	3.67
CC	4.22	3.17	3.46
CD _{WP,GW}	5.27	3.71	4.16
CD _{EB}	4.68	3.11	3.72
CD _{CC}	4.50	2.86	3.50

Table 3: Performance of different color constancy methods on the linear GreyBall dataset (lower is better).

method	mean (°)	median (°)	trimean (°)
do nothing	15.62	14.00	14.56
Low-level statistics-based methods			
GW	13.01	10.96	11.53
WP	12.68	10.50	11.25
SoG	11.55	9.70	10.23
general GW	11.55	9.70	10.23
GE1	10.58	8.84	9.18
GE2	10.68	9.02	9.40
Learning-based methods			
PG	11.79	8.88	9.97
EG	12.78	10.88	11.38
IG	11.81	8.93	10.00
HVLI	9.73	7.71	8.17
NIS	9.87	7.65	8.29
EB	7.97	6.46	6.77
CC	8.73	7.07	7.43
CD _{WP,GW}	10.27	7.33	8.20
CD _{EB}	8.46	5.63	6.73
CD _{CC}	8.79	5.66	6.89

below 3°, which was experimentally shown to be an acceptable error for human observers (Finlayson et al., 2005) (Fredembach and Finlayson, 2008), the results of CD_{WP,GW} are in most cases satisfying. Fig. 6 shows the results of chromatic adaptation based on illumination estimation of several methods on the NUS datasets. Beyond the fact that the CD_{WP,GW}, it also demonstrates how even very simple statistics of a realistic illumination distribution can significantly improve the initially low accuracy of simple methods up to the level to outperform state-of-the-art methods in most cases.

For the GreyBall dataset and its linear version CD_{WP,GW} outperformed all statistics-based methods and compared very well with the learning-based ones being less accurate than only a small number of them. The other tested combination CD_{EB} outperformed all other methods in terms of median angular error proving that already accurate methods can have their accuracy improved even further in a simple way.

3.4 Computational speed

Since the computation cost of voting is negligible, the computational cost the Color Dog depends only on the combined computational costs of its voters. According to the tests performed in (Cheng et al., 2014b) White-patch and Gray-world method are the two fastest methods and even their combined computation time is low, so CD_{WP,GW} is not only accurate, but also a fast method suitable for implementation in digital cameras with limited computational power.

4 CONCLUSIONS AND FUTURE RESEARCH

A new global illumination estimation learning-based methods has been proposed. It uses other methods' illumination estimations to vote for the most appropriate illumination from a predefined set with almost no additional computational cost. The method demonstrates how even some of the least accurate methods can be improved up to the level of outperforming most of the other methods while at the same time keeping their advantages. It also demonstrates that the accuracy of already very accurate methods can be improved even further. In future some more sophisticated and accurate voting methods should be researched.

ACKNOWLEDGEMENTS

This research has been partially supported by the European Union from the European Regional Development Fund by the project IPA2007/HR/16IPO/001-040514 "VISTA - Computer Vision Innovations for Safe Traffic."

REFERENCES

- A. Gijsenij, T. G. and van de Weijer, J. (2014). Color Constancy — Research Website on Illuminant

- Estimation.
- 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. (2015). Color Cat: Remembering Colors for Illumination Estimation. *Signal Processing Letters, IEEE*, 22(6):651–655.
- Buchsbaum, G. (1980). A spatial processor model for object colour perception. *Journal of The Franklin Institute*, 310(1):1–26.
- Cardei, V. C., Funt, B., and Barnard, K. (2002). Estimating the scene illumination chromaticity by using a neural network. *Journal of the Optical Society of America A*, 19(12):2374–2386.
- 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., Prasad, D., and Brown, M. S. (2014a). On Illuminant Detection.
- Cheng, D., Prasad, D. K., and Brown, M. S. (2014b). Illuminant estimation for color constancy: why spatial-domain methods work and the role of the color distribution. *Journal of the Optical Society of America A*, 31(5):1049–1058.
- 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.
- Deng, Z., Gijssenij, A., and Zhang, J. (2011). Source camera identification using Auto-White Balance approximation. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 57–64. IEEE.
- Ebner, M. (2007). *Color Constancy*. The Wiley-IS&T Series in Imaging Science and Technology. Wiley.
- Finlayson, G. D., Hordley, S. D., and Morovic, P. (2005). Colour constancy using the chromatic constraint. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 1079–1086. IEEE.
- 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.
- Fredembach, C. and Finlayson, G. (2008). Bright chromagenic algorithm for illuminant estimation. *Journal of Imaging Science and Technology*, 52(4):40906–1.
- 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.
- Gehler, P. V., Rother, C., Blake, A., Minka, T., and Sharp, T. (2008). Bayesian color constancy revisited. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE.
- Gijssenij, A. and Gevers, T. (2007). Color Constancy using Natural Image Statistics. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8.
- Gijssenij, A., Gevers, T., and Van De Weijer, J. (2011). Computational color constancy: Survey and experiments. *Image Processing, IEEE Transactions on*, 20(9):2475–2489.
- Gijssenij, A., Gevers, T., and Van De Weijer, J. (2012). Improving color constancy by photometric edge weighting. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(5):918–929.
- 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.
- Japkowicz, N. and Shah, M. (2011). *Evaluating Learning Algorithms*. Cambridge University Press.
- Joze, H. R. V. and Drew, M. S. (2012). Exemplar-Based Colour Constancy. In *British Machine Vision Conference*, pages 1–12.
- Joze, H. R. V., Drew, M. S., Finlayson, G. D., and Rey, P. A. T. (2012). The Role of Bright Pixels in Illumination Estimation. In *Color and Imaging Conference*, volume 2012, pages 41–46. Society for Imaging Science and Technology.

- L. Shi, B. F. (2014). Re-processed Version of the Gehler Color Constancy Dataset of 568 Images.
- 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.
- Van De Weijer, J., Gevers, T., and Gijsenij, A. (2007a). Edge-based color constancy. *Image Processing, IEEE Transactions on*, 16(9):2207–2214.
- Van De Weijer, J., Schmid, C., and Verbeek, J. (2007b). Using high-level visual information for color constancy. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE.
- Vassilvitskii, S. and University, S. (2007). *K-means: Algorithms, Analyses, Experiments*. Stanford University.
- Zhang, B. and Batur, A. (2014). Illumination Estimation Using Natural Scene Statistics. US Patent App. 14/188,670.