Color Beaver: Bounding Illumination Estimations for Higher Accuracy

Karlo Koščević, 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, Genetic Algorithm, Illumination Estimation, Image Enhancement, White Balancing.

Abstract: The image processing pipeline of most contemporary digital cameras performs illumination estimation in order to remove the influence of illumination on image scene colors. In this paper an experiment is described that examines some of the basic properties of illumination estimation methods for several Canon's camera models. Based on the obtained observations, an extension to any illumination estimation method is proposed that under certain conditions alters the results of the underlying method. It is shown that with statistics-based methods as underlying methods the proposed extension can outperform camera's illumination estimation in terms of accuracy. This effectively demonstrates that statistics-based methods can still be successfully used for illumination estimation in digital cameras. The experimental results are presented and discussed. The source code is available at https://ipg.fer.hr/ipg/resources/color_constancy.

1 INTRODUCTION

Among many abilities human visual system (HVS) can recognize colors of objects regardless of scene illumination. This ability is known as color constancy (Ebner, 2007). Achieving computational color constancy is an important pre-processing step in image processing pipeline as different scene illuminations may cause the image colors to differ as shown in figure 1. In order to remove the influence of illumination color, accurate illumination estimation followed by chromatic adaptation must be preformed. For both tasks the following image **f** formation model, which includes Lambertian assumption, is most often used:

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

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 c-th color channel. With the assumption of uniform illumination the problem can be simplified, as now $\mathbf{x}$ is removed from $I(\lambda, \mathbf{x})$ and the observed light source color is given as:

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

(a) (b)

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

For a successful chromatic adaptation, what is required is only the direction of **e** (Barnard et al., 2002). Since it is very common that only image pixel values **f** are given and both $I(\lambda)$ and $\rho(\lambda)$ remain unknown, calculating **e** is an ill-posed problem. To solve this problem, additional assumptions must be made, which leads to many color constancy methods that are divided into two major groups. First group of methods are low-level statistic-based methods like White-patch (Land, 1977; Funt and Shi, 2010), 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), Gray-Edge (1st and 2nd order) (Van De Weijer et al., 2007a), using bright and dark colors (Cheng et al., 2014). The second group is formed of learning-based methods like gamut mapping (pixel, edge, and intersection based) (Finlayson

Koščević, K., Banić, N. and Lončarić, S.
Color Beaver: Bounding Illumination Estimations for Higher Accuracy.
DOI: 10.5220/0007393701830190
In *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2019)*, pages 183-190
ISBN: 978-989-758-354-4

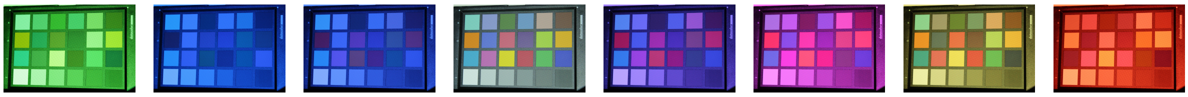

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 2: Color checker cast with projector light of various colors.

et al., 2006), using high-level visual information (Van De Weijer et al., 2007b), natural image statistics (Gijsenij and Gevers, 2007), Bayesian learning (Gehler et al., 2008), spatio-spectral learning (maximum likelihood estimate, and with gen. prior) (Chakrabarti et al., 2012), simplifying the illumination solution space (Banić and Lončarić, 2015a; Banić and Lončarić, 2015b; Banić and Lončarić, 2015b), using color/edge moments (Finlayson, 2013), using regression trees with simple features from color distribution statistics (Cheng et al., 2015), performing various kinds of spatial localizations (Barron, 2015; Barron and Tsai, 2017), using convolutional neural networks (Bianco et al., 2015; Shi et al., 2016; Hu et al., 2017; Qiu et al., 2018).

While learning-based method have a much higher accuracy, it are low-level statistics-based methods that are still being widely used among digital camera manufacturers since they are much faster and often more hardware-friendly than learning-based methods. This is also one of the reasons why statistics-based methods are still important for research. Nevertheless, since cameras are commercial systems, the fact that they still widely use statistics-based methods is not publicly stated. In this paper an experiment is described that examines some of the basic properties of illumination estimation methods for several Canon's camera models. Based on the obtained observations, an extension to any illumination estimation method is proposed that under certain conditions alters the results of the underlying method by bounding them to a previously learned region in the chromaticity plane. The bounding procedure is simple and does not add any significant cost to the overall computation. It is shown that with statistics-based methods as underlying methods the proposed extension can outperform camera's built-in illumination estimation in terms of accuracy. This effectively demonstrates that statistics-based methods can still be successfully used for illumination estimation in digital cameras' pipelines.

The paper is structured as follows: Section 2 lays out the motivation for the paper, in Section 3 the proposed method is described, Section 4 shows the experimental results, and Section 5 concludes the paper.

2 MOTIVATION

2.1 Do statistics-based Methods Matter?

Digital cameras are being used ever more widely, especially with the growing number of smartphones. This definitely means that the results of the research on computational color constancy now also have a higher impact so the importance of this research grows, especially when considering that it is an ill-posed problem. In literature and in the reviews of papers it is sometimes claimed that there is little purpose in researching low-level statistics-based methods since there are now much more accurate learning-based methods that significantly outperform them in accuracy. In contrast to that many experts with experience in the industry claim that many commercial white balancing systems are still based on low-level statistics-based methods. The main reason for that is their simplicity, low cost of implementation, and hardware-friendliness. If this is indeed so, then the research on such methods is definitely still important and should be further conducted and supported.

To check to what degree all these claims are true, it should be enough to examine some of the white balancing systems widely used in commercial cameras. In the world of professional cameras Canon has been the market leader for 15 years (Canon, 2018) and in 2018 it held an estimated 49% of the market share (PhotoRumors, 2018). Since practically every digital camera performs white balancing in its image processing pipeline, it can be claimed that Canon's white balancing system is one of the most widely spread ones. However, since Canon is a commercial company, full details of the white balancing system used in its digital cameras are not publicly known.

2.2 Learning from Existing Systems

One approach to gain more information on Canon's white balancing system is to look at the results of illumination estimation for various images taken under illumination of numerous colors. The following three camera models have been used to perform this experiment: EOS 550D, EOS 6D, and EOS 750D.

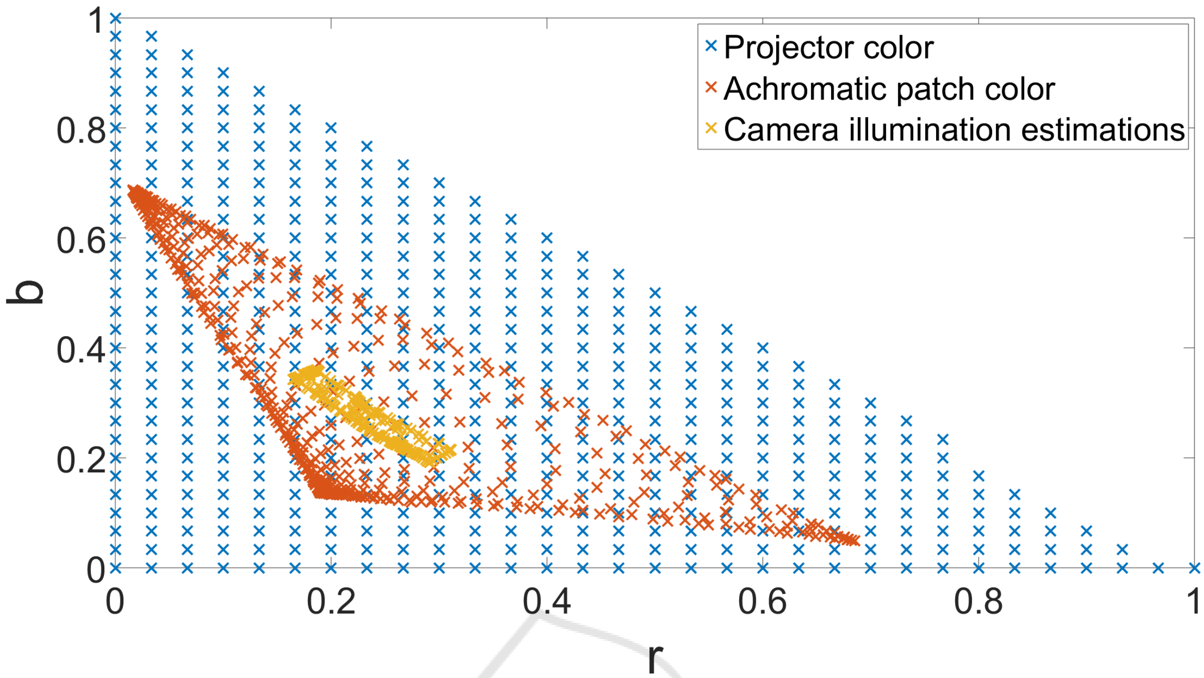

Figure 3: Comparison of chromaticities of projector light color, color of the second achromatic color checker patch, and camera's illumination estimation for Canon EOS 550D in the *rb*-chromaticity plane. The red chromaticity is shown on the *x* axis, while the blue chromaticity is shown on the *y* axis.

The experiment was conducted in a dark room where only a projector has been used as a light source. The projector was used to cast illumination of various colors, with chromaticities evenly spread in the chromaticity plane, on a color checker as shown in Figure 2. These images of the color checker were taken with every of the three previously mentioned cameras.

Although the illumination color was supposed to be computationally determined by projecting specifically created content, due to the projector and camera sensor characteristics the effective illumination color is altered. Its value as perceived by the camera can be read from the achromatic patches in the last row of the color checker and it serves as the ground-truth illumination for the given image. Ideally, it is this color that an illumination estimation method should predict.

Finally, the last step of the experiment was to check the results of illumination estimation performed by each of the cameras. The results of a camera's illumination estimation for a taken image can be reconstructed from the Exif metadata stored in the image file. The fields needed for this are *Red Balance* and *Blue Balance*, which have the values of channel gains i.e. the factors by which the red and blue channels have to be multiplied to perform chromatic adaptation. For practical reasons in cameras the green gain is fixed to 1. The combined inverse values of these gains give the illumination estimation vector. When this vector is normalized, it represents the chromaticity of camera's illumination estimation, which can be directly used to calculate the estimation accuracy by comparing it to the ground-truth illumination.

A comparison between the chromaticities for projected illumination color, achromatic patch color, and camera illumination estimations for Canon EOS 550D camera is given in Figure 3. The values read from achromatic white patches are squeezed with respect to the ones sent by the projector, but a more interesting observation is that none of the camera's illumination estimation are outside of a surface that resembles a parallelogram. As shown in Figure 4, similar results are obtained for other used camera models as well. Although there are some differences between the parallelograms mostly visible on two opposite sides, the parallelograms otherwise mostly cover a similar space in the chromaticity plane.

2.3 Observations

Based on these observations it can be concluded that one of the core properties of Canon's white balancing system is limiting its illumination estimation so that they do not appear outside of a polygon very similar to a parallelogram. Such limitation can be justified by the goal of avoiding unlikely illuminations and thus minimizing the occurrence of too high errors. This can be useful if it can be assumed that the expected illuminations are similar to black body radiation, but sometimes it can be an disadvantage if artificially colored illumination sources are present like in Figure 2.

On the other hand, there is little that can be said about the white balancing system's behavior inside of the parallelogram. Nevertheless, the limitation observation is already useful because of its potential to limit maximum errors for illumination estimations. As for the behavior of illumination estimation inside the parallelogram, a possible solution is to use some of the already existing methods. Additionally, it can be immediately remarked that a parallelogram is a relatively regular quadrangle and polygon in general.

At least two questions can be raised here: first, is there a better quadrangle i.e. polygon for bounding the illuminations, and second, which method to use as the baseline underlying method that gets bounded?

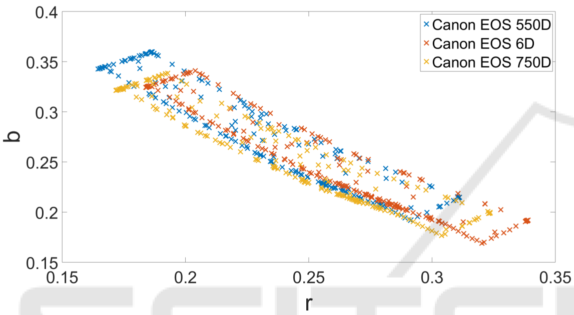

Figure 4: Comparison of cameras' illumination estimation for Canon EOS 550D, Canon EOS 6D, and Canon EOS 750D. The red chromaticity is shown on the x axis, while the blue chromaticity is shown on the y axis.

3 PROPOSED METHOD

Inspired by the bounds used by Canon cameras observed in Figure 4 and in order to give an answer to the two questions from the previous section, in this paper a new method i.e. extension to any chosen underlying illumination estimation method is proposed. The extension learns a bounding polygon with an arbitrary number of vertices that is used to restrict the illumination estimations of the initially chosen underlying method to the chromaticity region specified by the bounding polygon. As explained in the previous section, the motivation for this are the observations of boundaries used by Canon cameras and it can be applied to any illumination estimation method.

A genetic algorithm is used to learn the boundaries. First, the illumination estimations for the initially chosen underlying method are calculated on a given set of images. The boundary polygon population of size s is initialized by taking randomly chosen ground-truth illumination chromaticities as polygon vertices. Empirically, it has been shown that the four-point polygons i.e. quadrangles are generally a good fit for illumination restriction and there is no significant gain when the number of points is increased. The fitness function calculation for a given quadrangle is based on the ground-truth illuminations and the restricted illuminations that are the result of applying the boundary polygon to the underlying method's illumination estimations. Empirically, it has been concluded that the negative sum of the median angular error and a tenth of the maximum angular error is generally a good fitness function; angular error is explained in more detail in Section 4.1. More formally, if $\mathbb{A} = \{a_1, \dots, a_n\}$ is the set of angular errors on n images, then the chosen fitness function is given as

$$\mathbf{f}(\mathbb{A}) = -\left(\operatorname{med}(\mathbb{A}) + \frac{1}{10}\max(\mathbb{A})\right). \quad (3)$$

The maximum error was also included in the fitness function in order to discourage quadrangles that perform very well on the majority of the images, but have poor performance of several outliers. As the selection method the 3-way tournament selection (Mitchell, 1998) with random sampling is used. Averaging crossover function of the two of the best individuals produces a new child which is randomly mutated. The quadrangle with the lowest fitness value among the three ones chosen in the selection procedure is replaced in the current population by the newly created child quadrangle. The mutation is done by translating each vertex of a bounding polygon by the value from the normal distribution with $\mu = 0$ and $\sigma = 0.2$. Mutation rate, which states whether the whole individual should me mutated, is set to 0.3. After all training iterations have finished, the boundary quadrangle with the highest fitness value is chosen as the final result. Figure 5 shows an example of a learned quadrangle.

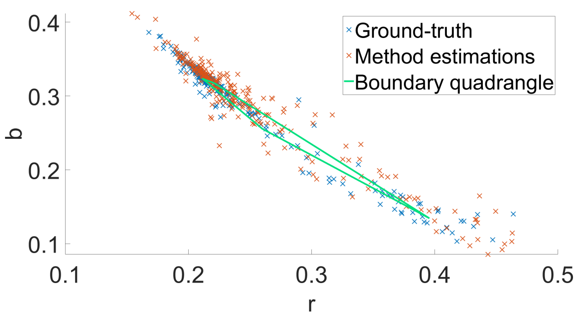

Figure 5: Example of a learned boundary quadrangle for the Canon1 dataset (Cheng et al., 2014) in the chromaticity plane. The red chromaticity is shown on the x axis, while the blue chromaticity is shown on the y axis.

Since the proposed extension bounds illumination estimations and beavers are known to bound water flows by building dams, the proposed extension was named Color Beaver. In the rest of the paper extending a method **M** by the Color Beaver extension will be denoted as Color Beaver + **M**. The training procedure for Color Beaver is summarized in Algorithm 2.

Algorithm 1: Color Beaver Training.

Input: training images $\mathbb{I}$, ground truth $\mathbb{G}$, method $\mathbf{M}$, iterations number N, population size s, fitness function $\mathbf{f}$
Output: boundary polygon $\mathbf{P}$

$\mathbb{E} = estimateIllumination(\mathbb{I}, \mathbf{M})$
$\mathbb{P} = initializePolygonPopulation(s)$
for $\mathbf{i} \in \{1, ..N\}$ **do**
  $\mathbf{t_1}, \mathbf{t_2}, \mathbf{t_3} = tournamentSelection(\mathbb{P}, 3, \mathbf{f})$
  $\mathbf{t}' = crossover(\mathbf{t_1}, \mathbf{t_2})$
  $\mathbf{t}'.mutateMaybe(0.3)$d
  $\mathbb{R} = restrictIllumination(\mathbb{E}, \mathbf{t}')$
  $\mathbb{P}.ReplaceExistingWith(\mathbf{t_3}, \mathbf{t}')$
end for
$\mathbf{P} = \mathbb{P}.GetFittest(\mathbf{f})$

Algorithm 2: Color Beaver Application.

Input: image $\mathbf{I}$, method $\mathbf{M}$, boundary polygon $\mathbf{P}$
Output: illumination estimation $\mathbf{e}$

$\mathbf{e}_M = estimateIllumination(\mathbf{I}, \mathbf{M})$
$\mathbf{e} = restrictIllumination(\mathbf{e}_M, \mathbf{P})$

4 EXPERIMENTAL RESULTS

4.1 Experimental Setup

Eight linear NUS datasets (Cheng et al., 2014) and the Cube dataset (Banić and Lončarić, 2017) have been used to test the proposed extension and compare its performance to the one of other methods. All these datasets have linear images, which is also expected by the model described by Eq. (3). The ColorChecker dataset (Gehler et al., 2008; Shi and Funt, 2018) has not been used because of much confusion that is still present in many papers due to of its misuses in the past (Lynch et al., 2013; Finlayson et al., 2017).

The most commonly used accuracy measure among many proposed (Gijsenij et al., 2009; Finlayson and Zakizadeh, 2014; Banić and Lončarić, 2015a) is the angular error. It is the angle between the vectors of illumination estimation and the ground-truth illumination. When the angular errors obtained on each individual image of a given benchmark dataset need to be summarized, one of the most important statistics is the median angular error (Hordley and Finlayson, 2004). Despite that fact, the geometric mean of several statistics including the median angular error has increasingly been gaining popularity in recent publications (Barron and Tsai, 2017) and the same format as there is also used in this paper.

For both the NUS datasets and the Cube dataset a three-fold cross-validation with folds of equal size was used like in previous publications. The source code for recreating the results reported later in the paper is publicly available at https://ipg.fer.hr/ipg/resources/color_constancy.

Table 1: Performance of different color constancy methods on the Cube dataset (Banić and Lončarić, 2017) in terms of angular error statistics (lower Avg. is better). The used format is the same as in (Barron and Tsai, 2017).

Algorithm	**Mean**	**Med.**	**Tri.**	**Best 25%**	**Worst 25%**	**Avg.**
White-Patch (Funt and Shi, 2010)	6.58	4.48	5.27	1.18	15.23	4.88
Gray-world (Buchsbaum, 1980)	3.75	2.91	3.15	0.69	8.18	2.87
Camera built-in	2.96	2.56	2.64	0.82	5.79	2.49
Color Tiger (Banić and Lončarić, 2017)	2.94	2.59	2.66	0.61	5.88	2.35
Shades-of-Gray (Finlayson and Trezzi, 2004)	2.58	1.79	1.95	0.38	6.19	1.84
2nd-order Gray-Edge (Van De Weijer et al., 2007a)	2.49	1.60	1.80	0.49	6.00	1.84
1st-order Gray-Edge (Van De Weijer et al., 2007a)	2.45	1.58	1.81	0.48	5.89	1.81
General Gray-World (Barnard et al., 2002)	2.50	1.61	1.79	0.37	6.23	1.76
Color Beaver Camera + built-in (proposed)	1.70	0.96	1.15	0.31	4.38	1.20
Color Beaver + WP (proposed)	1.59	0.87	1.04	0.25	4.15	1.08
Restricted Color Tiger (Banić and Lončarić, 2017)	1.64	0.82	1.05	0.24	4.37	1.08
Color Dog (Banić and Lončarić, 2015b)	1.50	0.81	0.99	0.27	3.86	1.05
Smart Color Cat (Banić and Lončarić, 2015b)	1.49	0.88	1.06	0.24	3.75	1.04
Color Beaver + SoG (proposed)	1.51	0.81	1.00	0.22	3.97	1.01
Color Beaver + GW (proposed)	1.48	0.76	0.98	0.21	3.90	0.98

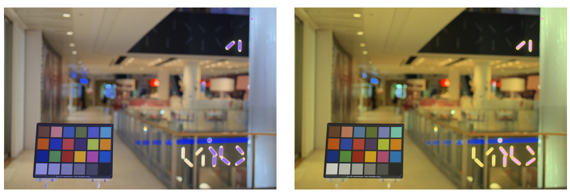

(a) (b)

Figure 6: A failure case for Color Beaver + SoG with chromatic adaptation results based on a) the restricted illumination estimation with angular error of 10.74° and b) the ground-truth illumination.

Table 2: Combined performance of different color constancy methods on eight NUS dataset in terms of angular error statisrics (lower Avg. is better). The used format is the same as in (Barron and Tsai, 2017).

Algorithm	Mean	Med.	Tri.	Best 25%	Worst 25%	Avg.
White-Patch (Funt and Shi, 2010)	9.91	7.44	8.78	1.44	21.27	7.24
Pixels-based Gamut (Gijsenij et al., 2010)	5.27	4.26	4.45	1.28	11.16	4.27
Grey-world (Buchsbaum, 1980)	4.59	3.46	3.81	1.16	9.85	3.70
Edge-based Gamut (Gijsenij et al., 2010)	4.40	3.30	3.45	0.99	9.83	3.45
Color Beaver + WP (proposed)	5.40	2.12	2.75	0.58	16.08	3.12
Shades-of-Gray (Finlayson and Trezzi, 2004)	3.67	2.94	3.03	0.98	7.75	3.01
Color Beaver + GW (proposed)	3.73	2.65	2.90	0.72	8.55	2.82
Natural Image Statistics (Gijsenij and Gevers, 2011)	3.45	2.88	2.95	0.83	7.18	2.81
Local Surface Reflectance Statistics (Gao et al., 2014)	3.45	2.51	2.70	0.98	7.32	2.79
2nd-order Gray-Edge (Van De Weijer et al., 2007a)	3.36	2.70	2.80	0.89	7.14	2.76
1st-order Gray-Edge (Van De Weijer et al., 2007a)	3.35	2.58	2.76	0.79	7.18	2.67
Bayesian (Gehler et al., 2008)	3.50	2.36	2.57	0.78	8.02	2.66
General Gray-World (Barnard et al., 2002)	3.20	2.56	2.68	0.85	6.68	2.63
Spatio-spectral Statistics (Chakrabarti et al., 2012)	3.06	2.58	2.74	0.87	6.17	2.59
Bright-and-dark Colors PCA (Cheng et al., 2014)	2.93	2.33	2.42	0.78	6.13	2.40
Corrected-Moment (Finlayson, 2013)	2.95	2.05	2.16	0.59	6.89	2.21
Color Beaver + SoG (proposed)	2.86	1.99	2.21	0.59	6.62	2.17
Color Tiger (Banić and Lončarić, 2017)	2.96	1.70	1.97	0.53	7.50	2.09
Color Dog (Banić and Lončarić, 2015b)	2.83	1.77	2.03	0.48	7.04	2.03
Shi et al. 2016 (Shi et al., 2016)	2.24	1.46	1.68	0.48	6.08	1.74
CCC (Barron, 2015)	2.38	1.48	1.69	0.45	5.85	1.74
Cheng 2015 (Cheng et al., 2015)	2.18	1.48	1.64	0.46	5.03	1.65
FFCC (Barron and Tsai, 2017)	1.99	1.31	1.43	0.35	4.75	1.44

4.2 Accuracy

Tables 1 and 2 show the comparisons between the accuracies of methods extended by the proposed extension and other illumination estimation methods. It can be seen that all of the extended methods outperform their initial non-extended versions. As a matter of fact, the extended version of the Shades-of-Gray method outperforms the camera built-in method. Additionally, the extended versions also outperform many learning-based methods. All these results demonstrate the usability of the proposed extension. An example of a failure case for the proposed extension of Shades-of-Gray is shown in Figure 6.

While other methods such as Gray-edge could also have been tested and shown in the Tables, Shades-of-Gray was already good enough to outperform camera's built-in methods. Extending Gray-edge also increases its accuracy, but Gray-edge is slower than Shades-of-Gray (Cheng et al., 2014), more complex, and it requires additional memory. Hence it was left out of the testing procedures since Shades-of-Gray is already sufficient to successfully answer the questions that were raised in this paper.

4.3 Discussion

The fact that statistics-based methods extended by the proposed method outperform camera built-in illumination estimation methods is significant for drawing further conclusions about the nature of camera's illumination estimation methods. Namely, if extended statistics-based methods outperform them, it can be freely stated that statistics-based are good enough to be used in digital cameras. Additionally, it may be that the extended method managed to outperform the camera's built-in methods because that they are also statistics-based, which in turn confirms that cameras do indeed use such method. In any of these two cases it can be concluded that research on statistics-based methods still has a large field of applications and the obtained results only further prove its importance.

5 CONCLUSIONS

An experiment was conducted to examine some of the details of built-in illumination estimation methods for several Canon camera models. Inspired by the observed results, an extension to any underlying illumination estimation method has been proposed. It limits the values of the illumination estimations of the underlying method by forcing it to stay inside a previously learned region in the chromaticity plane without adding any significant computation cost. By limiting some of the best-known statistics-based methods, the obtained accuracy outperforms the one of cameras' built-in methods. This effectively demonstrates that by only using slightly modified statistics-based methods it is possible to be more accurate than contemporary cameras. It also proves the claim that statistics-based methods can and probably are used for illumination estimation in digital cameras. Future research will include looking for new method modifications that result in even higher estimation accuracy.

ACKNOWLEDGEMENTS

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 Workshopn (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.

Canon (2018). Canon celebrates 15th consecutive year of No.1 share of global interchangeable-lens digital camera market.

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. 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.

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.

Gao, S., Han, W., Yang, K., Li, C., and Li, Y. (2014). Efficient color constancy with local surface reflectance statistics. In *European Conference on Computer Vision*, pages 158–173. Springer.

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.

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

Gijsenij, A. and Gevers, T. (2011). Color constancy using natural image statistics and scene semantics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):687–698.

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. (2010). Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision*, 86(2):127–139.

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.

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.

Mitchell, M. (1998). *An introduction to genetic algorithms*. MIT press.

PhotoRumors (2018). 2018 Canon, Nikon and Sony market share (latest Nikkei, BCN and CIPA reports).

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, L. and Funt, B. (2018). Re-processed Version of the Gehler Color Constancy Dataset of 568 Images.

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. (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.