Computational Aesthetics in Graphics, Visualization, and Imaging (2009)
O. Deussen and P. Hall (Editors)

Consistent Scene Illumination using a Chromatic Flash

Min H. Kim[†] and Jan Kautz[‡]

Department of Computer Science, University College London


(a) Ambient Light Only (b) Ambient + Ordinary Flash (c) Ambient + Chromatic Flash

Figure 1: *Image (a) is a photograph of a sunset scene using ambient illumination only, lacking detail in dark areas. Image (b) is captured with an ordinary flash unit, changing the atmosphere of the scene illumination to an overly bluish color (note the illumination mismatch on the sphere and the legs). Image (c) is the image captured using our adaptive chromatic flash producing consistent scene illumination.*

Abstract

Flash photography is commonly used in low-light conditions to prevent noise and blurring artifacts. However, flash photography commonly leads to a mismatch between scene illumination and flash illumination, due to the bluish light that flashes emit. Not only does this change the atmosphere of the original scene illumination, it also makes it difficult to perform white balancing because of the illumination differences. Professional photographers sometimes apply colored gel filters to the flashes in order to match the color temperature. While effective, this is impractical for the casual photographer. We propose a simple but powerful method to automatically match the correlated color temperature of the auxiliary flash light with that of scene illuminations allowing for well-lit photographs while maintaining the atmosphere of the scene. Our technique consists of two main components. We first estimate the correlated color temperature of the scene, e.g., during image preview. We then adjust the color temperature of the flash to the scene's correlated color temperature, which we achieve by placing a small trichromatic LCD in front of the flash. We demonstrate the effectiveness of this approach with a variety of examples.

1. Introduction

The human visual system has the ability to adapt to different illumination conditions and to perceive surface colors independently of the illumination, called *chromatic adaptation* [dL86]. In digital photography this is simulated using *white balancing* through *computational color constancy* [Hun06], ensuring that white walls appear white. However, this proves difficult if the scene is illuminated by multiple light sources with different color temperatures. This is especially true with flash photography where the flash casts a bluish light due to its xenon lamp. This mismatch in illumination often makes objects stand out unnaturally, see Figure 1(b) for instance. Simulating color constancy is difficult in this case, as the illumination varies spatially and ordinary white balancing methods sacrifice either of the illumination conditions [HMP*08].

[†] e-mail: m.kim@cs.ucl.ac.uk

[‡] e-mail: j.kautz@cs.ucl.ac.uk

We propose a simple yet powerful technique to solve this illumination mismatch in flash photography without requiring any manual intervention. The main idea is to estimate the correlated color temperature of the scene, and then to adjust the flash's correlated color temperature according to that of the scene, which we achieve by mounting a small LCD screen in front of an ordinary flash. This enables the automatic capture of low-light scenes in a single shot without altering the atmosphere of the scene illumination.

Our main contributions include:

- A highly efficient method to estimate the correlated color temperature of scene illumination.
- A complete pipeline to reduce the mismatch of correlated color temperatures between scene and auxiliary illumination.

2. Background and Related Work

Correlated Color Temperature *Color temperature* is defined as the spectral power distribution of a Planckian blackbody radiator [WS82]. Even though many real-world illuminants are not exactly equal to any of the chromaticities of a blackbody radiator, we can compute the *correlated color temperature* (CCT) [HK75], which refers to the closest matching temperature. In our work, we estimate the CCT of a scene, in order to match it with the flash unit. While this assumes the scene illumination to be on the blackbody locus, it acts as a constraint which allows us to find rather good estimates.

Color Constancy Techniques for estimating the correlated color temperature are usually a part of *computational color constancy* [dL86], which simulate the human visual system's chromatic adaptation in digital imaging. Conceptually, these algorithms first *estimate* the correlated color temperature and then *balance* the white-point of the image accordingly. In the context of our work, we use color constancy in two ways. First, we propose an efficient method to estimate the correlated color temperature of a scene in order to match the flash illumination; and second, we apply chromatic adaptation, CIECAT02 [CIE04] for final display (for simplicity, in-camera white balancing can also be used).

Many color constancy methods have been proposed and we can only mention the most related ones; for a more complete overview, see Hordley [Hor06]. In order to estimate the unknown scene illumination from camera signals only, assumptions are usually made about aspects of real-world images. The gray-world method [Buc80, vdWG05] assumes that the average reflectance or color derivative in a scene is gray, whereas the maxRGB method [Lan77] assumes the respectively brightest channel levels in an image correspond to the white point. Instead, prior information about the gamut distribution can be acquired in a learning phase, which is used in the color-by-correlation method for instance [FHH01]. Statistical prior probability of the training data set can be used to improve the performance of the gray-world method [BCF02, GG07, GRB*08]. This requires a large set of training data and long precomputation times.

Despite the large variety of available methods, no algorithm can be regarded as universal. In practice, the gray-world and maxRGB approaches perform well on natural, real-world images [Hor06, GG07]. We therefore use an enhanced version of the gray-world algorithm to estimate the scene's CCT, which is inspired by Barnard et al.'s method [BCF02]. However, we derive a linear transform from real-world training images with radiometric measurements instead of synthetic images, and we further apply a weighting scheme that combines the maxRGB and gray-world methods.

Hardware Design The illumination mismatch problem has been addressed by several camera manufacturers – in the form of patents. Iwasaki [Iwa98] proposed a flash light sensing structure to do white balancing according to the reflected flash light. Adjustable (colored) flash bulbs have been proposed using LEDs, as well as using mixtures of two or more gases (xenon bulbs) [ZH89, PG90, AHM02]. Even the combination of flash and LCD screen has been patented [BBS04]. In contrast to these patents, we actively demonstrate the benefits of an adaptive chromatic flash. Furthermore, rather than targeting a specific hardware design for consumer-level cameras, we focus on the efficient computation of correlated color temperatures from camera signals as well as the precise matching of CCTs.

Spatially Varying White Balancing Ebner [Ebn04] and Hsu et al. [HMP*08] have proposed spatially varying white balancing techniques with similar aims to our work, i.e., handling images with two different illuminants, such as flash and ambient illumination. While the latter technique works quite well, it requires manual intervention after the image has been taken, whereas the former achieves good results in a few cases. In contrast, we avoid spatially varying illuminants in the first place by matching scene and flash temperature, which allows photographers to immediately check the quality of the resulting image.

Flash/No-flash Pairs Eisemann and Durand [ED04] propose a technique to merge pairs of flash/no-flash images in order to keep the warm atmosphere of the no-flash image with the detail of the flash image. Agrawal et al. [ARNL05] have a similar goal but base their method on gradient coherence and projection. Yuan et al. [YSQS07] deconvolve no-flash images in order to remove blur from camera shake, where the kernel is derived from a flash/no-flash pair. Our goal is similar, we want to keep the atmosphere of the no-flash image, but increase detail using a flash. However, we take the approach of changing the flash chromaticity so that only a single flash picture needs to be taken.

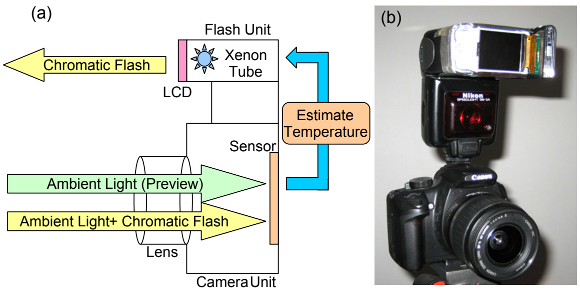

Figure 2: *(a) Schematic overview of our method. During image preview, the correlated color temperature of the ambient illumination is estimated. The color of the LCD, which is attached to the flash, is adjusted in order to produce the same color temperature as the ambient illumination. Finally, the scene is captured by the camera using the chromatic flash maintaining the characteristics of the ambient illumination. Our prototype is shown in (b).*

3. Adaptive Chromatic Flash

Our approach consists of two main steps. First, the color temperature of the scene illuminant needs to be estimated, e.g., during image preview. Second, we illuminate the scene with a chromatic flash so that its correlated color temperature matches that of the scene. We have implemented a hardware prototype of the proposed method, see Figure 2, with which all the results were generated.

3.1. Estimating the Scene Illumination

The camera signal ρ (for each color channel $k = r, g, b$) is the integral of the product of surface reflectance $S(\lambda)$, camera response function $D_k(\lambda)$ (e.g., influenced by color filters), and irradiance $E(\lambda)$ over all wavelengths λ:

$$\rho_k = \int E(\lambda) S(\lambda) D_k(\lambda) \, d\lambda \, . \tag{1}$$

We characterize $D_k(\lambda)$ [BF02], which allows us to get (linearized) estimates of the radiant power $\Phi = E(\lambda)S(\lambda)$.

However, both $E(\lambda)$ and $S(\lambda)$ are unknown. In order to adjust our chromatic flash, we need to estimate the correlated color temperature T of the *scene illuminant* $E(\lambda)$. We start from the gray-world assumption that the average of all surface reflectances in a scene is a neutral reflectance [Buc80]. However, as mentioned in [BCF02, GG07, GRB*08], real-world statistical data shows that the average is different from perfect neutral reflectance. Unlike previous database-based gray-world methods [BCF02, GG07, RML03] that either use synthetic training images or training images without knowing the actual scene illuminant, we use a database of real-world photographs as well as accurately measured scene illuminants $E(\lambda)$.

We first captured 35 training images of real-world scenes (see Figure 3) under different illumination conditions with a color temperature T_m ranging from 2000K to 7500K, which we measured on a Spectralon tile that was placed in each scene using a Jeti Specbos 1200 spectroradiometer. The Spectralon tile was always oriented such that it was facing the main light source. It was usually removed from the scene when the training images were photographed (see Figure 3).

The radiant power value Φ of each pixel (in each image) are then projected onto the blackbody locus using Holm and Krochmann's method [HK75], which is also used by the spectroradiometer to estimate the CCTs of the training data, yielding the (per pixel) correlated color temperature T:

$$\arg\min_T \left[(u_e - u_T)^2 + (v_e - v_T)^2 \right]^{1/2} , \tag{2}$$

where (u_e, v_e) are the radiance chromaticity coordinates of the pixel (derived from their radiance value) and T is the temperature of the nearest point (u_T, v_T) on the Planckian locus. The color temperatures T_i of pixels P_i within each image are then averaged together using a weighted average (similar to gray-world):

$$\bar{T} = \frac{\sum_i T_i \cdot w(P_i)}{\sum_i w(P_i)} , \tag{3}$$

Our weighting function $w()$ is proportional to the luminance of a pixel, i.e., zero weights are applied to the pixels with smallest luminance and a weight of one is applied to the brightest pixels.

From this training data, we then derive a simple affine transformation $T_m = a \cdot \bar{T} + b$ that maps from $\bar{T}$ to the accurately measured T_m. We estimate the two parameters a and b of this model using linear regression:

$$\mathbf{M}_T = (\bar{\mathbf{T}}^t \bar{\mathbf{T}})^{-1} \bar{\mathbf{T}}^t \mathbf{T}_m \, . \tag{4}$$


(a) 3480 K
(b) 4029 K
(c) 5432 K
(d) 5432 K
(e) 5759 K
(f) 6105 K
(g) 6620 K
(h) 6660 K
(i) 7485 K

Figure 3: *Examples of the training images. We use raw sensor signals (discarding the camera's auto white balance) and the spectral power distribution of the scene illumination (measured on a Spectralon tile) as our training data.*

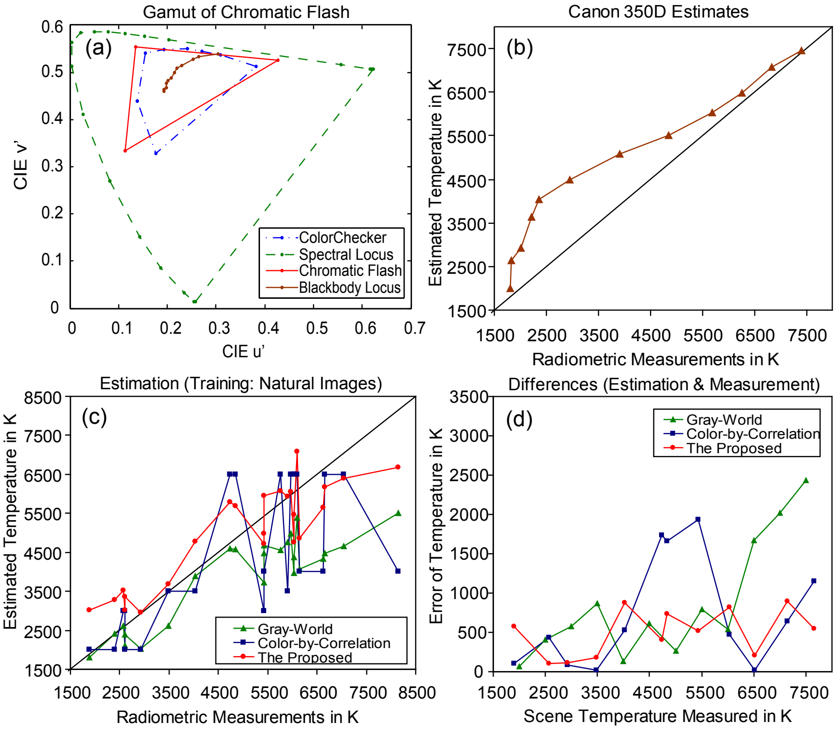

Figure 4: *(a) Gamut coverage of our chromatic flash. The LCD with xenon flash covers the range of blackbody radiation. (b) Correlated color temperature estimates from the Canon 350D. (c) Result of temperature estimation using the training data of natural images (all 35). (d) Difference between temperature estimation and radiometric measurement of new test images.*

where $\bar{\mathbf{T}}$ refers to the vector containing all training CCTs $\bar{T}$ and $\mathbf{T}_m$ refers to the vector containing all measured CCTs T_m, and $\mathbf{M}_T$ is a matrix containing the two parameters. For any new image, we simply compute $\bar{T}$ and map to the actual color temperature T_a with $\mathbf{M}_T$.

Traditional gray-world methods average trichromatic primaries first and then compute the correlated color temperature from the average. However, we have found that first computing color temperatures and building a weighted average of those yields better results (correlation coefficient of $R^2 = 0.86$ vs. $R^2 = 0.79$).

Initially, we experimented with training images of a GretagMacbeth DC chart instead of natural images. While their average color temperatures $\bar{T}$ were highly correlated with the measured color temperatures T_m ($R^2 = 0.99$), the derived linear transform did not generalize well to natural images.

Alternatively, one might want to use simple photodetectors to estimate the scene illuminant. However, this would not yield as good results, as important information about the scene is lost (e.g., the colors of bright areas). We therefore make use of the additional information contained in the images.

3.2. Chromatic Flash Illumination

Our flash prototype consists of an ordinary flash with a small RGB LCD panel mounted in front, see Figure 2. After estimating the color temperature T_a of the scene, we set the color of the LCD panel so that the flash emits light with a color temperature of T_a. To this end, we calibrate our chromatic flash in order to set the color temperature as accurately as possible.

The exposure time of the xenon flash bulb is less than a few milliseconds, which cannot be measured by our spectroradiometer. Instead, we calibrate the color of the chromatic flash with an accurately characterized camera [KK08]. We emit flashes with each of the LCD's red, green, and blue primaries turned on separately (different intensity levels) and capture physically meaningful XYZ images for each of them [KK08]. The different intensity levels are used to calibrate the LCD's gamma parameters γ_k (non-linear curve fitting). The brightest red, green, and blue XYZ images are used to derive a linear transform $\mathbf{M}_F$ that maps from the LCD's RGB primaries to XYZ. Given a CCT T and its associated chromaticity φ, we look up the corresponding XYZ coordinates (table according to [WS82]) and use $\mathbf{M}_F^{-1}$ to map into the LCD's primaries (plus gamma correction). We choose the brightest combination of LCD primaries for the given CCT in order to maximize transmittance.

While an RGB panel is sufficient to produce all possible CCTs rather accurately, the chromatic flash does not reproduce the exact spectrum of the original scene illuminant. In theory this may lead to an object's appearance being slightly off [WHD03], but we have not noticed any artifacts in our experiments.

4. Results

All our results were captured with a Canon 350D digital camera and a Nikon SB-24 flash unit with an attached LCD screen. Figure 4(a) presents the gamut boundary of our chromatic flash, compared with GretagMacbeth ColorChecker. The blackbody radiation falls into the gamut of chromatic flash, which means the whole range of color temperature of the blackbody can be synthesized.

Figure 4(c) and (d) demonstrate that our database-based gray-world algorithm estimates the color temperature rather accurately for both the training as well as new natural images (see throughout the paper and additional material). These results compare favorably to the original gray-world and color-by-correlation methods.

Figure 5(a) and (b) demonstrate the difference maps of correlated color temperatures of each pixel (closeup of a sphere in Figure 1). Image (a) shows the difference map of CCTs of each pixel between the ambient scene and the ordinary flash-lit scene. The center of the sphere in image (a) shows significant differences of color temperatures. Image (b) shows the difference map between the ambient and the chromatic flash-lit scene. The front area of the sphere in image (b) is illuminated by our chromatic flash. There is virtually no difference in CCTs (less than 300K), demonstrating the effectiveness of our approach.

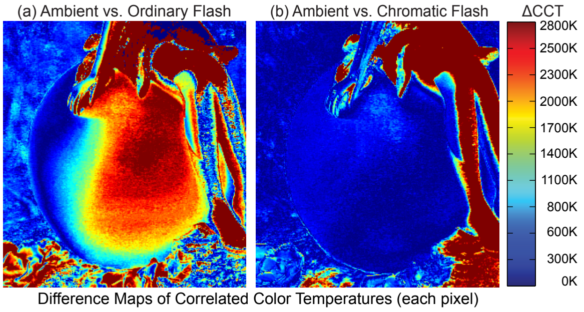

Figure 5: *Image (a) presents a difference map of correlated color temperatures between the ambient only scene and the ordinary flash-lit scene (closeup of the sphere from Figure 1). Image (b) shows a difference map of CCTs between the ambient-only scene and the chromatic flash-lit scene. Red areas indicate high temperature differences and blue areas indicate almost identical color temperatures. As can be seen in (b), the proposed chromatic flash provides auxiliary illumination without altering the original color temperature of the scene. Note that a change of color temperature is expected in areas that were originally in shadow, as they now receive light from the chromatic flash; e.g., see sphere and legs in (b).*

Figure 6 shows results from of our chromatic flash. Image (a) is taken with ambient lighting only – parts of the statue are fully shadowed. Image (b) is taken with an additional ordinary flash and shows a mismatch in illumination. Image (c) and (d) are both captured by the proposed chromatic flash but are displayed using different white balancing (WB) methods (camera WB and CIECAT02 [CIE04]). Both images avoid the mismatch of ambient and auxiliary flash illumination, producing better colors. Image (e) and (f) show two manual white-balancing alternatives (instead of auto white-balancing) for the ordinary flash image. However, even manual adjustments cannot reach the illumination consistency of our method.

Figure 7 shows the results of our chromatic flash prototype with different natural images. As can be seen, the images are illuminated consistently with our method and objects do not stand out unnaturally. Furthermore, this enables us to use traditional white balancing methods (two different ones are shown in (c) and (d)) as opposed to spatially varying white balancing, which requires manual intervention [HMP*08].

4.1. Discussion

In many cases, our color temperature estimation method is more accurate than the original gray-world or gamut-based model, even though we only used 35 training images. Yet, it allows us to compute an estimate in milliseconds (or even less when only a subset of pixels is used), which is necessary for lag-free photography. Of course, when an image deviates too much from our training data, the color temperature estimate is less accurate. Note that the camera's internal temperature estimate may not be directly applicable, as it is often skewed to accommodate user preferences. For instance, the brown sigmoidal curve in Figure 4(b) shows the Canon 350D's color temperature estimation of the GretagMacbeth training images (derived from white balancing multipliers), which indicates a deliberate choice to overestimate the color temperature.

When the scene illuminant moves far from the locus, the performance of our algorithm will degrade as we assume the illuminant to lie on the locus. However, in our experience, this case does not seem to occur frequently in natural scenes. Of course, extreme cases such as tinted light bulbs will be difficult to handle for our method. In addition, our method seems to perform well, even if the new images are not well represented in our training database. For instance, there is no similar training image to the color chart example in Figure 7 and the example from Figure 6.

Pictorial Comparison between Ordinary Flash & Our Method (Chromatic Flash)

WB Variation on Ordinary Flash

Figure 6: *Image (a) is a photograph of a statue under the gas street lamp using ambient illumination only, lacking detail on the statue. Image (b) is captured with a flash unit, showing the mismatch of the ambient and auxiliary flash light – the statue stands out bluish from the warm environment. Images (c) and (d) is captured with the proposed chromatic flash (using different white-balancing algorithms for reference). Our chromatic flash (c and d) produces a consistent appearance and no visible mismatch in illumination. For completeness, images (e) and (f) show alternative white-balancing for the ordinary flash image: ambient for (e), and flash for (f). This results in a yellowish tint on the left of the statue for image (e) and overly bluish for image (f).*

Although our method provides an effective solution to color temperature problems in flash photography, it inherits the limitations of ordinary flash units, e.g., reflections on glossy surfaces. The LCD panel does take away quite some light which makes a higher transmittance panel or powerful flash a better choice for this method.

5. Conclusion

We have presented a method that improves flash photography in terms of color reproduction by matching the flash's color temperature to that of the scene. Our proposed technique preserves the aesthetics of the original scene's ambient illumination and enables the use of traditional white balancing methods. Considering the small cost of an LCD display, the method can easily be integrated into cameras with flash units.

References

[AHM02] Amerson F. C., Hubel P. M., Motta R. J.: *Auxiliary Illuminating Device Having Adjustable Color Temperature*. US Patent 6,379,022, 2002.

[ARNL05] Agrawal A., Raskar R., Nayar S. K., Li Y.: Removing photography artifacts using gradient projection and flash-exposure sampling. *ACM Trans. Graph. (Proc. SIGGRAPH) 24*, 3 (2005), 828–835.

[BBS04] Bloom D. M., Bianchi M. J., Sobol R. E.: *Auxiliary Illuminating Device Having an Adjustable Color Temperature by Controlling the Amount of Light Passing through Color Filters*. US Patent 6,755,555, 2004.

[BCF02] Barnard K., Cardei V., Funt B. V.: A comparison of computational color constancy algorithms – part I: Methodology and experiments with synthesized data. *IEEE Trans. Image Processing 11*, 9 (2002), 972–984.

[BF02] Barnard K., Funt B. V.: Camera characterization for color research. *Color Res. Appl. 27*, 3 (2002), 153–164.

[Buc80] Buchsbaum G.: A Spatial Processor Model for Object Colour Perception. *J. the Franklin Institute 310*, 1 (1980), 1–26.

[CIE04] CIE: *TC8-01 Technical Report, A Colour Apperance Model for Color Management System: CIECAM02*. Publication CIE 159-2004, CIE, 2004.

[dL86] d'Zmura M., Lennie P.: Mechanisms of color constancy. *J. Opt. Soc. Am. 3* (1986), 1662–1672.

[Ebn04] Ebner M.: Color constancy using local color shifts. In *Proc. ECCV* (2004), pp. 276–287.

[ED04] Eisemann E., Durand F.: Flash photography enhancement via intrinsic relighting. *ACM Trans. Graph. (Proc. SIGGRAPH) 23*, 3 (2004), 673–678.

[FHH01] Finlayson G. D., Hordley S. D., Hubel P. M.: Color by correlation: A simple, unifying framework for color constancy. *IEEE Trans. Pattern Anal. Mach. Intel. 23*, 11 (2001), 1209–1221.

[GG07] Gijsenij A., Gevers T.: Color constancy using natural image statistics. In *Proc. IEEE CVPR* (2007), pp. 1–8.

[GRB*08] Gehler P. V., Rother C., Blake A., Minka T.,

(a) Ambient Light Only (b) Ambient + Ordinary Flash (c) Ambient + Chromatic Flash (AutoWB) (d) Ambient+Chromatic Flash(CIECAT02)

Figure 7: *Results of our chromatic flash. Column (a) presents the captured image with ambient illumination only. Column (b) provides the results when using an ordinary flash unit with white balancing set to mixed illumination (ambient and flash). Flash illuminated areas appear relatively bluish. Column (c) is the result of our proposed method with camera auto white balancing, which preserves the atmosphere of original scene illumination. Column (d) is also the result of our proposed method but with manual white balancing (CIECAT02).*

Sharp T.: Bayesian color constancy revisited. In *Proc. IEEE CVPR* (2008), pp. 1–8.

[HK75] Holm W., Krochmann J.: Ermittlung der ähnlichsten Farbtemperatur. *Die Farbe 24*, 1/6 (1975), 91–96.

[HMP*08] Hsu E., Mertens T., Paris S., Avidan S., Durand F.: Light Mixture Estimation for Spatially Varying White Balance. *ACM Trans. Graph. (Proc. SIGGRAPH) 27*, 3 (2008), 1–7.

[Hor06] Hordley S. D.: Scene illuminant estimation: Past, present, and future. *Color Res. Appl. 31*, 4 (2006), 303–314.

[Hun06] Hung P.-C.: Color theory and its application to digital still cameras. In *Image Sensors and Signal Processing for Digital Still Cameras*, Nakamura J., (Ed.). 2006, pp. 205–221.

[Iwa98] Iwasaki H.: *Chromatic Balancer for Flash Cameras*. US Patent 5,749,002, 1998.

[KK08] Kim M. H., Kautz J.: Characterization for high dynamic range imaging. *Computer Graphics Forum (Proc. EUROGRAPHICS) 27*, 2 (2008), 691–697.

[Lan77] Land E. H.: The retinex theory of color vision. *Scientific American 237*, 6 (1977), 108–128.

[PG90] Petrakos S., Gaewsky J. P.: *Variable Color-Output Strobe*. US Patent 4,978,892, 1990.

[RML03] Rosenberg C. R., Minka T. P., Ladsariya A.: Bayesian color constancy with non-gaussian models. In *NIPS* (2003), MIT Press.

[vdWG05] van de Weijer J., Gevers T.: Color constancy based on the grey-edge hypothesis. In *Proc. ICIP* (2005), pp. 722–725.

[WHD03] Wenger A., Hawkins T., Debevec P.: Optimizing color matching in a lighting reproduction system for complex subject and illuminant spectra. In *Proc. EGSR* (2003), pp. 249–259.

[WS82] Wyszecki G., Stiles W. S.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, 1982.

[YSQS07] Yuan L., Sun J., Quan L., Shum H.-Y.: Image deblurring with blurred/noisy image pairs. *ACM Trans. Graph. (Proc. SIGGRAPH) 26*, 3 (2007), 1:1–1:10.

[ZH89] Zeltner U., Hauser H. P.: *Flash Appartus with Color Temperature Control*. US Patent 4,853,600, 1989.