

Flash Lighting Space Sampling

Matteo Dellepiane, Marco Callieri, Massimiliano Corsini, Paolo Cignoni,
and Roberto Scopigno

Visual Computing Lab, ISTI-CNR, Pisa, Italy

`name.surname@isti.cnr.it`

`http://vcg.isti.cnr.it`

Abstract. Flash light of digital cameras is a very useful way to picture scenes with low quality illumination. Nevertheless, especially low-end cameras integrated flash lights are considered as not reliable for high quality images, due to known artifacts (sharp shadows, highlights, uneven lighting) generated in images. Moreover, a mathematical model of this kind of light seems difficult to create. In this paper we present a color correction space which, given some information about the geometry of the pictured scene, is able to provide a space-dependent correction of each pixel of the image. The correction space can be calculated once in a lifetime using a quite fast acquisition procedure; after 3D spatial calibration, obtained color correction function can be applied to every image where flash is the dominant illuminant. The correction space presents several advantages: it is independent from the kind of light used (provided that it is bound to the camera), it gives the possibility to correct only determinate artifacts (for example color deviation) introduced by flash light, and it has a wide range of possible applications, from image enhancement to material color estimation.

Keywords: Color, shading, shadowing, and texture.

1 Introduction

The vast diffusion of digital cameras started a revolution in the way photography is used to measure the appearance of real-world objects. However, despite the significant improvements of this kind of devices, there are several issues, regarding the correctness of the acquired data, that need a careful evaluation. Particularly, the setup and calibration of the lighting environment is often one of the most limiting aspects for simplified appearance acquisition techniques.

All cameras are equipped with a flash: this tool can provide a practical, easy, cheap controlled lighting environment. Unfortunately, built-in flashes usually produce a variety of undesirable effects, like uneven lighting (overexposed in the near field and dark in the distance), highlights and distracting sharp shadows. Moreover, "white balance" setting for flash is not effective because it applies the same correction throughout the image. More interesting results could be obtained by knowing the geometry of the scene, and using a model of the behavior of flash light. However, a mathematical modeling of flash light spatial behavior can be

hard to be obtained, because of the large variety of camera/flash models, and the irregular spatial light distribution produced by the flash reflectors and lenses.

For these reasons we propose Flash Lighting Space Sampling (FLiSS), a correction space where a correction matrix is associated to each point in the camera field of view. Once that basic information about the geometry of the scene is known, the proposed structure permits to correct each pixel according to the position of the corresponding point in space. This structure proves to be simple and effective, and it has several advantages, which will be presented and discussed during the analysis of results and the concluding remarks.

2 Related Work

The work proposed in this paper shares some aspects with various subjects, such as computational photography, image enhancement and lighting modeling and estimation. Some of the most interesting and related works will be presented.

Flash/No-Flash Digital Photography. The use of flash/no-flash pairs to enhance the appearance of photographs is a relatively recent research topic where several interesting works appeared. The *continuous flash* [17] was a seminal work, where flash and no-flash pairs were combined to create adjustable images. Two almost contemporaneous papers [9,23] proposed techniques to enhance details and reduce noise in ambient images, by using flash/no-flash pairs. These works provide features for detail transfer, color and noise correction, shadows and high-lights removal. While the systems are not completely automatic, very interesting results can be easily obtained. The goal of a more recent work [1] is to enhance flash photography by introducing a flash imaging model and a gradient projection scheme, to reduce the visual effects of noise. Flash/no-flash pairs are used by [21] to detect and remove ambient shadows.

Color constancy and white balance. White balance is a key issue in the context of the *color constancy problem*, that studies the constancy of perceived colors of surfaces under changes in the intensity and spectral composition of the illumination. Several works in this field rely on the assumption that a single illuminant is present: the enhancement of photos can be based on geometric models of color spaces [10], statistical analysis of lights and colors [11] or natural images [12], study of the edges of the image [25].

Another group of papers deals with mixed lighting conditions. Methods can be semi-automatic [20] or automatic. Automatic methods usually work well under quite strong assumptions, like hard shadows and black-body radiators lights [19] or localized gray-world model [8]. A very recent work [18] proposes a white balance technique which renders visually pleasing images by recovering a set of dominant material colors using the technique proposed by [22]. One of the assumptions is that no more than two light types (specified by the user) illuminate the scene. Most of the cited works share some of the main hypotheses of our

method. Nevertheless, the knowledge of some information about the geometry of the scene eliminates the need for other restricting assumptions (such as smooth illumination, gray-world theory, need of user interaction).

Illumination estimation and light models. The works in the field of illumination estimation have two principal aims: the estimation of the lighting of an environment [7,24] or the measure of the characteristics of a luminary. Our work is more related to the second group.

One of the first attempts to model both the distant and the near behavior of a light source is the *near-field photometry* approach of Ashdown [2]. Near-field photometry regards the acquisition of a luminary by positioning a number of pinhole cameras (or moving a single camera) around it, and measuring the incident irradiance on a CCD sensor. The results are mapped onto an hemicycle that represents the final model of the luminary acquired. Heidrich et al used a similar method [16] by moving the camera on a virtual plane and representing the light sources with a Lumigraph [14]. This representation was named *canned light source*. More recently, Goesele et al. [13] improved the near-field photometric approach using a correction filter to compensate the fact that a digital camera is not a real pinhole camera. Our approach recalls near-field photometry; the main difference is that we estimate the data to “correct” the effect of the light source on known colors, instead of building a model of the light source of interest.

3 Definition of FLiSS

The aim of our work is to build a spatial color correction function that associates a specific color correction procedure to each point in the camera frustum space. We call this particular data structure *color correction space*. Such an approach allows to override the limitations assumed in most of the color correction approaches [3,4], that is that the illumination is constant, or easily model-able, across the scene. Our main assumptions are: flash light can be considered the dominant illumination in the scene; the light interaction can be described using just sRGB space (we do not account full spectra data); surfaces are non-emitting. These hypotheses are common among existing techniques which deal with single illumination, and they cover most of the real cases. Typically, the color calibration of digital photographs consists in taking a snapshot of a pre-defined color reference target, such as a Macbeth ColorCheckerTM or an AGFA IT8TM, placed near the subject of interest, and estimating the parameters of a transformation that maps the colors of the reference target in the image into its real colors.

A quite simple to model the correction is a linear affine transformation $c' = Ac + B$. Obviously, due to the nonlinear nature of image color formation, this kind of correction is a rough approximation and many other approaches could be used [3,4]. Moreover, the correction is effective for the image parts that are close (and with a similar illumination) to the reference target. On the other hand, in practice, this simple and compact approach works reasonably well in most cases.

The linear transformation can be written as a 4×3 matrix:

$$\begin{pmatrix} R' \\ G' \\ B' \end{pmatrix} = \begin{bmatrix} c_{11} & c_{12} & c_{13} & c_{14} \\ c_{21} & c_{22} & c_{23} & c_{24} \\ c_{31} & c_{32} & c_{33} & c_{34} \end{bmatrix} \begin{pmatrix} R \\ G \\ B \\ 1 \end{pmatrix} \quad (1)$$

In the following we refer to this matrix as the color correction matrix \mathcal{C} and to its elements as the correction parameters. We explain how we estimate \mathcal{C} in Section 4.

Roughly speaking, the parameters of \mathcal{C} have the following meaning: (c_{11}, c_{22}, c_{33}) are related to the change in contrast of the color; $(c_{12}, c_{13}, c_{21}, c_{23}, c_{31}, c_{32})$ are related to the color deviation caused by the color of the flash light (if the flash is purely white light, these components tend to zero); (c_{14}, c_{24}, c_{34}) are related to the intensity offset. We use the term contrast in the sense that the multiplication for the coefficients expands the range of values of the channels.

Given the assumptions above, we can finally define FLiSS. Flash Lighting Space Sampling is a color correction space where a color correction matrix is associated to each point in a camera frustum. The correction space will be calculated starting from several sampled points in the camera space. The process of correcting an image will, for each pixel in the image, use the appropriate correction matrix according its corresponding position in the camera space. Due the continuous nature of the correction space, the correction will prove to be reliable even without a precise digital model of the scene, so that an approximate reconstruction such as the ones generated, for example, by stereo matching could be used as well.

4 Acquisition Procedure and Data Processing

The aim of our work was to try to build a procedure which could be used in a general case, without using prototypal or expensive devices. The computation of our color correction space is necessary only once (or very few times) in a camera lifetime. But even with this assumption, it was important to define some as simple and fast as possible acquisition procedures. Hence, we decided to sample the camera space view frustum by taking flash lighted photos of a small color target, calculating the correction matrix in those points and subsequently building the entire space by interpolation. We performed the acquisition with three different models of digital cameras, shown in Figure 1(left). These models are representative of three categories of non-professional cameras: compact, digital SLR (single-lens reflex) and digital SLR with external flash.

As a color target we used a Mini Macbeth ColorChecker; its size (about $3.5'' \times 2.5''$) allows the assumption that the light variation across it is negligible. To sample the view frustum of the camera we sliced it with several planes of acquisition at different distances, moving the Mini Macbeth in different positions for each plane. We divided a distance range between 50 and 220 cm in 7 planes, with 25 positions for each plane, as shown in Figure 1(right). The color




Fig. 1. Left: Digital cameras used for light space sampling. Right: the scheme of acquisition for flash space sampling, with a snapshot of the acquisition setup.

target was placed on a tripod, and always faced towards the camera. For each position multiple snapshots were taken, in order to deal with the known variability of flash behavior, and keeping a fixed exposure time and aperture during the entire procedure.

The snapshots were acquired in sRGB RAW format to avoid any other internal processing by the digital camera, except for the Casio compact camera, with which we were forced to use JPEG images. The acquisition procedure of all the needed images for each model took a couple of hours. The processing of acquired data was subdivided in two main phases. In the first one we calibrated all the acquired images, using the color target reference. In the second phase we built the color correction space through parameters interpolation. Figure 2 shows a schematization of the entire data processing.

Color Calibration

The color calibration was done using the calibration model explained in Section 3. The parameters of the matrix C were estimated by solving a linear system with the c_{ij} as unknowns. Our parameter estimation algorithm takes inspiration from the RANSAC approach.

Since four colors could be sufficient to write a system with 12 equations and 12 unknowns, several random combinations of colors are used to find the best solution in terms of quality. The quality of the color correction is evaluated




Fig. 2. Data processing. Left: The acquired color targets are calibrated generating a set of color correction matrices, and then interpolated on each plane. Right: Each correction parameter is interpolated to fill the whole camera space.

considering the CIELab distance [6] between the real colors and the corrected ones. Hence, after extracting the Mini Macbeth in a semi-automatic way and segmenting each color patch, several permutations of the colors are tried in order to find the best estimation in terms of quality. This robust approach produces better results from a perceptual point of view than a standard least square approach.

After calibration, each image is associated to: the 12 parameters of the matrix \mathcal{C} , the position of the Mini Macbeth relative to the camera view and the distance at which the image has been taken. Since several shots of the same position of the target are taken, obtained values are the means of all the obtained estimations. Before this, the data are further processed with a statistical analysis [15] to remove the outliers from the acquired data.

Data Interpolation

Starting from a set of color correction matrices for several points in the space, we could try either to fit a mathematical function or to calculate the intermediate color correction matrix through interpolation. Here, we opt for interpolation, leaving the first approach as an interesting direction of future research. Given a point p in camera space, we calculate the corresponding color correction matrix as the linear interpolation in the squared distance between the camera center O and the point p in the following way:

$$\mathcal{C}_{ij}(p) = \mathcal{C}_{ij}(p_1) + (d^2 - d_1^2) \frac{\mathcal{C}_{ij}(p_2) - \mathcal{C}_{ij}(p_1)}{d_2^2 - d_1^2} \quad (2)$$

where $\mathcal{C}(x)$ indicates the color correction matrix at the point x ; p_1 and p_2 are the intersection points between the line starting from O and passing through p . These points lie on the acquisition plane immediately before and after p (Figure 2right); d_1 , d and d_2 are the distances between the point O and p_1 , p and p_2 .

Since $\mathcal{C}(p_1)$ and $\mathcal{C}(p_2)$ are not known in advance, they have to be estimated to evaluate (2). In fact, only few positions on the acquisition plane are measured: hence another interpolation is required. For this planar interpolation we use radial basis method, with gaussians centered on each sample: standard deviation σ defines how much each sample influences the neighbors. In formulas:

$$\mathcal{C}_{ij}(p) = \frac{\sum_{i=1}^N \mathcal{C}_{ij}(P_i) \exp \left[-\frac{(p - P_i)^2}{\sigma^2} \right]}{\sum_{i=1}^N \exp \left[-\frac{(p - P_i)^2}{\sigma^2} \right]} \quad (3)$$

where N is the total number of samples for the plane and P_i are the positions of samples.

In conclusion, in order to calculate a color correction matrix for a point in the camera space, we first need to calculate two linear interpolations in the acquisition planes, then a linear squared interpolation for the final result.

Regarding the practical implementation, the pre-computed interpolation planes are stored as floating point textures plus, for each texture, some additional information (the distance from the camera center and a scale factor). With this representation, the correction algorithm can be entirely implemented on the GPU, and the pre-computation of the interpolation planes reduces the evaluation of $C_{ij}(p_1)$ and $C_{ij}(p_2)$ in (2) to two texture lookups.

5 Data Analysis

The analysis of the data obtained from the acquisition was performed before the validation of results. A first analysis was on the value ranges obtained for the single coefficients of the matrix. Table 1 shows the statistics relative to all the coefficients calculated for the Nikon camera. The single mean values of

Table 1. Statistics of single coefficients for Nikon camera

	Contrast			Intensity offset			Color deviation					
Coefficient	c_{11}	c_{22}	c_{33}	c_{14}	c_{24}	c_{34}	c_{12}	c_{13}	c_{21}	c_{23}	c_{31}	c_{32}
Mean value	1.20	1.25	1.33	52.07	50.90	50.47	.07	-.06	.005	-.12	.07	-.31
Variance	0.06	0.09	0.09	81.57	123.5	141.3	.001	.003	.002	.009	.001	.012
Min value	0.83	0.84	0.87	27.5	15.4	14.2	-.06	-.43	-.29	-.67	-.21	-.85
Max value	2.59	2.84	2.86	72.0	75.0	76.8	.23	.01	.08	.01	.12	-.11

each group are very similar, and variance describes a general stability of data. Moreover, color deviation coefficients describe the flash as very near to a white light (with a slight deviation in green channel). Contrast and Intensity offset groups show quite low values in variance, but, as it could be expected, the needed modification of color values increases a lot with the distance. Further information about the properties of flash light can be inferred from the analysis of the azimuthal and normal sections of the correction space. In Figure 3 the planes associated to coefficient c_{11} are shown: isolines help in understanding the shape of the light. The light wavefront is quite similar for all the models, but several differences arise as well. The most regular profile is the one associated to the external flash (Canon camera): this is probably the most reliable kind of illumination. Nevertheless, the isolines show that the shape of the light is not similar to a sphere, but it could be better approximated with series of lobes. Moreover, it can be noted that the maximum value of the coefficient is higher respect to the other two examples: this is because the external flash was set with the largest possible field of view, resulting in a very bright illumination for near points, with the need of more correction for the far ones. The Casio associated space is the less regular one, and this is probably due to the fact that the camera model was the least expensive one, and the images were processed and stored directly in JPEG format. Anyway, an interesting observation is that the correction space seems to be slightly shifted on the right, like the position of the flash with respect to the lens (Figure 1).




Fig. 3. Azimuthal and normal plane for parameter c_{11} in camera space. (Left) Nikon camera. (Center) Canon camera. (Right) Casio camera.

6 Results

In order to evaluate the results of the correction introduced by FLiSS, we created a setup where the color values in selected regions were known in advance: the scene was formed by a series of identical objects, set at different distances from the camera. The reference objects were simple structures (formed by bricks of different colors) of LEGO DUPLO[®], whose size and color are known to be practically identical for all components. The reference RGB color value of each brick was calculated by calibration with the Mini Macbeth. Seven blocks of LEGO[®] bricks were put in a setting shown in Figure 4(left), and the corresponding 3D model of this configuration was created. This scene has the advantage that, in every single photos, identical elements are present in different parts of the image and at different distances in space. Images were aligned to the 3D model using a registration tool: thus color correction was performed on each pixel which was framing a structure.

One possible criticism to our work could be: would a simple model, for example a point light, be enough to provide a good correction? For this reason, we modeled the flash as a point light and we corrected the same images using this model. We obtained an acceptable estimation of the intensity of the light, by modeling the light degradation with a quadratic law and taking into account the knowledge of the distance of each pixel in the scene, and its reference color: we used this light estimation to correct the color value. Results of corrections on one of the images are shown in Figure 4(left): the original, the Fliss corrected,

and the point light corrected images are shown. While in both cases the colors across the scene seem similar, obtained color values are clearly different between the two results. In order to check the accuracy of color correction we measured the perceptual distance between produced colors and the previously calculated reference colors. We chose the CIE76 Delta E measure, which is the Euclidean distance between the CIELAB values associated to two colors. If Delta E is smaller than one, colors are perceived as identical; for values smaller than eighteen, the colors can be considered as very similar.

Figure 4 shows the results of the calculation of these distance values. The top boxes display reference colors, then for each piece (each line is associated to the piece number indicated in Figure 4 left) the three columns represent the average color value for original, Fliss corrected and point light corrected image. The Delta E value for each color shows the distance with respect to the reference. It is quite clear that, although producing similar colors, the point light correction returns results which are different respect to the reference: only red color is accurately corrected. Slight improvements are introduced only for distant objects. In order to achieve better results, probably a different modeling for each channel would be necessary. On the contrary, Fliss correction proves to be very reliable, with an average Delta E value which is always smaller than ten. Only in the case of block number 1 the correction is less effective: this can be due to the fact that the object is out of focus, or the color in the original image may be saturated.

A second validation experiment was performed in order to show that the correction introduced by Fliss can be reliable also with low quality geometric information of the scene. We reconstructed the geometry of a common scene starting from a few photos, using the Arc3D web service [26]. Figure 5 shows the

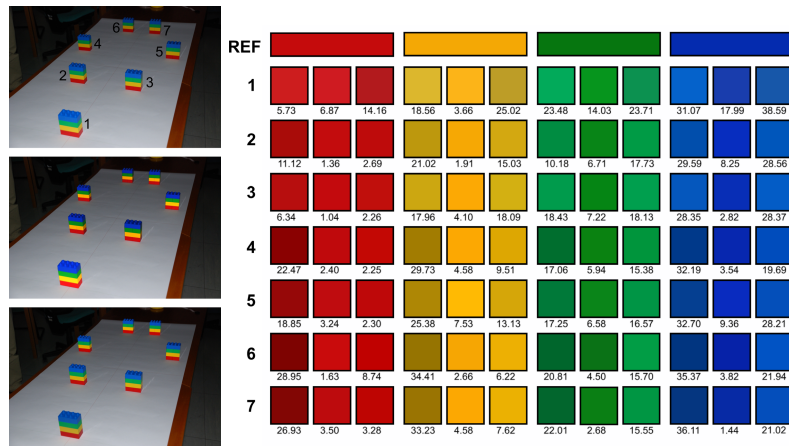


Fig. 4. Left: example of images used for validation: top, original; center, Fliss corrected; bottom, point light model corrected. Only the Lego blocks part of the image was corrected. Right: correction accuracy estimation: for each Lego block, average color and Delta E value (respect to reference) in original (left column), Fliss corrected (center column) and Point Light corrected (right column) images are shown




Fig. 5. Example of color correction using low quality geometry: Top left, original flash image; Top right, a snapshot of the extracted geometry; Bottom, Fliss corrected image.

starting image, together with the geometry obtained, and the corrected image (the correction was applied only where geometry data was present). While the obtained 3D model is only an approximated representation of the geometry in the image, the result of FLiSS use is satisfying: for example the color of the tablecloth after correction is the same throughout the scene.

Analysis of validation tests brings us to three main conclusions: first of all, Fliss is a very reliable way to correct images; moreover, simpler light models are not able to achieve comparable results. Finally, the approach is reliable regardless of the quality of the geometry associated to the image.

7 A Practical Application

To further show the potentiality of FLiSS, we tested its impact in the framework of a photographic mapping tool for 3D scanned models. The tool follows the approach described in [5] where the color assigned to each vertex of the 3D model is computed as weighted sum of the contributions of all the photos which project on that vertex. These weights are calculated by considering several metrics and their use guarantees the better quality data to prevail, obtaining a smooth transition between photos, without loss of detail. This mapping approach is fast, robust and easy to be implemented, but the artifacts produced by flash, like highlights, hard shadows and non uniform illumination cannot be automatically corrected.



Fig. 6. Top: an example of color correction on a flash photograph: the original image (left); a normalized map of the offset correction on the image (center); the corrected image (right). Bottom: a colored 3D model, visualized with no illumination (only color values assigned to vertices). Color correction OFF in left image, ON in right one.

We selected a test set of objects (colored Nativity statues of different heights and materials), to assess the quality and the impact of the flash light correction. We generated the 3D models of the statues with a Minolta 3D scanner and the photos with the Nikon D40x and flash light as the principal illuminant. Images were aligned on the models, and each image of the set was corrected using FLiSS: an example of color correction on a single image is shown in Figure 6.

With an accurate estimation of light position (obtained with a simple procedure involving a simple LEGO[©] calibration rig), it was also easy to detect the zones where shadows and highlights were present, and remove them by using the redundancy of data provided by the images of the set.

The effects of color correction on a final 3D Model can be seen in Figure 6: the model on the right, obtained with color correction, appears much more "flat" with respect to the model produced without color correction and artifact removal (left). In conclusion, these images show one of the possible applications of our color correction space, where the effectiveness of results is strengthened by the fact that acquisition of color information was fast and easy.

8 Concluding Remarks and Future Work

We presented FLiSS, a novel structure to apply a spatial color correction on images taken using flash light. The method requires basic information of the geometry of the scene, and it's able to automatically correct the color values in the image. FLiSS needs to be estimated only once in a camera lifetime, and this calibration procedure is adequately fast and easy. FLiSS has several advantages: it is robust and flexible, and although its use best fits with the case of flash light, the structure can be used with any kind of light which is bound to the camera. Future work can include both improvements in the proposed structure and exploitation of its possible uses. Regarding the first issue, further effort could be put in trying to find even simpler acquisition procedures, or reduce the number of samples positions. Regarding possible future applications, FLiSS could be used in the context of image enhancement, by applying for example only part of the elements of the correction matrix to obtain visual pleasantness or enhanced readability. Moreover, by working on different flash-lighted images and extracting the local illumination data from the FLiSS, it should be possible to recover simple models of the optical properties of a surface. Finally, while the simple point light model for flash exhibits strong limitations, FLiSS data could be used to obtain a more complex model that could approximate in a very accurate way the behavior of the flash light.

Acknowledgements. This work was partially funded by EU IST IP 3DCOFORM and Tuscany Regional Project STAR-T. We would like to thank Michael Goesele for his kind support, Giuliano Kraft and Francesco Dellepiane for providing us the hardware for experimental studies.

References

1. Agrawal, A., Raskar, R., Nayar, S.K., Li, Y.: Removing photography artifacts using gradient projection and flash-exposure sampling. *ACM Tr. on Graphics* 24(3), 828–835 (2005)
2. Ashdown, I.: Near-Field Photometry: Measuring and Modeling Complex 3-D Light Sources. In: *ACM SIGGRAPH 1995 Course Notes*, pp. 1–15 (1995)
3. Barnard, K., Cardei, V., Funt, B.: A comparison of computational color constancy algorithms. I: Methodology and experiments with synthesized data. *Image Processing, IEEE Trans.* 11(9), 972–984 (2002)
4. Barnard, K., Martin, L., Coath, A., Funt, B.: A Comparison of Computational Color Constancy Algorithms Part II: Experiments With Image Data. *IEEE Trans. on Image processing* 11(9), 985 (2002)
5. Callieri, M., Cignoni, P., Corsini, M., Scopigno, R.: Masked photo blending: mapping dense photographic dataset on high-resolution 3d models. *Computer & Graphics* 32(4), 464–473 (2008)
6. Commission Internationale de l'Eclairage (CIE). *Colorimetry CIE 15* (2004)
7. Debevec, P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In: *SIGGRAPH 1998*, pp. 189–198. ACM Press, New York (1998)

8. Ebner, M.: Color constancy using local color shifts. In: Pajdla, T., Matas, J.(G.) (eds.) ECCV 2004. LNCS, vol. 3023, pp. 276–287. Springer, Heidelberg (2004)
9. Eisemann, E., Durand, F.: Flash photography enhancement via intrinsic relighting. In: ACM Trans. on Graphics, vol. 23, ACM Press, New York (2004)
10. Finlayson, G.D., Hordley, S.D.: Improving gamut mapping color constancy. IEEE Trans. on Image Processing 9(10), 1774–1783 (2000)
11. Finlayson, G.D., Hordley, S.D., Hubel, P.M.: Color by correlation: A simple, unifying framework for color constancy. IEEE Trans. on Pattern Analysis and Machine Intelligence 23(11), 1209–1221 (2001)
12. Gijsenij, A., Gevers, T.: Color constancy using natural image statistics. In: Int. Conf. on Comp. Vision and Pat. Recogn., Minneapolis, USA, June 2007, pp. 1–8 (2007)
13. Goesele, M., Granier, X., Heidrich, W., Seidel, H.-P.: Accurate light source acquisition and rendering. ACM Tr. Gr. 22(3), 621–630 (2003)
14. Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F.: The lumigraph. In: SIGGRAPH 1996, pp. 43–54 (1996)
15. Grubbs, F.: Procedures for detecting outlying observations in samples. Technometrics 11, 1–21 (1969)
16. Heidrich, W., Kautz, J., Slusallek, P., Seidel, H.-P.: Canned lightsources. In: Rend. Tech., pp. 293–300 (1998)
17. Hoppe, H., Toyama, K.: Continuous flash. Technical Report MSR-TR-2003-63, Microsoft Research (2003)
18. Hsu, E., Mertens, T., Paris, S., Avidan, S., Durand, F.: Light mixture estimation for spatially varying white balance. In: SIGGRAPH 2008, ACM Press, New York (2008)
19. Kawakami, R., Ikeuchi, K., Tan, R.T.: Consistent surface color for texturing large objects in outdoor scenes. In: ICCV 2005: Int. Conf. on Computer Vision, Washington, DC, USA, pp. 1200–1207 (2005)
20. Lischinski, D., Farbman, Z., Uyttendaele, M., Szeliski, R.: Interactive local adjustment of tonal values. ACM Tr. Gr. 25(3), 646–653 (2006)
21. Lu, C., Drew, M.S., Finlayson, G.D.: Shadow removal via flash/noflash illumination. In: W. on Mult. Signal Processing, pp. 198–201 (2006)
22. Omer, I., Werman, M.: Color lines: Image specific color representation. In: CVPR 2004, June 2004, vol. II, pp. 946–953. IEEE, Los Alamitos (2004)
23. Petschnigg, G., Szeliski, R., Agrawala, M., Cohen, M., Hoppe, H., Toyama, K.: Digital photography with flash and no-flash image pairs. In: SIGGRAPH 2004, pp. 664–672 (2004)
24. Unger, J., Wenger, A., Hawkins, T., Gardner, A., Debevec, P.: Capturing and rendering with incident light fields. In: EGRW 2003, pp. 141–149 (2003)
25. van de Weijer, J., Gevers, T.: Color constancy based on the grey-edge hypothesis. In: ICIP, pp. 722–725 (2005)
26. Vergauwen, M., Van Gool, L.: Web-based 3d reconstruction service. Mach. Vision Appl. 17(6), 411–426 (2006)