Towards the Estimation of Non-Uniform Illumination in Real-World Scenes

Christian Riess, Sven Pfaller, David Bernecker and Elli Angelopoulou
Pattern Recognition Lab
Friedrich-Alexander University Erlangen-Nuremberg
http://www5.cs.fau.de

Abstract

In a raw photographic image, the observed colors are a mixture of material reflectance and colored illumination. For further processing, the color influence of the light source is often considered a disturbance. In such cases, it is desireable to neutralize the color of the light source, often referred to as white balancing. Successful white balancing often requires an estimate of the color of the light source. Most existing illuminant color estimators assume homogeneous, globally uniform illumination. However, this assumption is often violated in real-world images.

In this work, we make an attempt towards estimating the colors and spatial distributions of multiple light sources in a scene. We estimate the color of the illumination locally, and define an objective function to select two or three representative illuminants from all local estimates. Although the objective function is relatively simple, qualitative results look promising.

1 Introduction

When capturing a scene with a camera, the observed colors are always a mixture of the material reflection and the illuminant spectrum. For several vision tasks like color-based object recognition, non-white illumination can be a serious source of error.

The goal of research in color constancy is to obtain a color-normalized version of the image. A common simplifying assumption is that an observed pixel color p_c is the product of material color m_c and illumination color i_c, i. e.

$$p_c = (m_c \cdot i_c)^\gamma \quad (c \in \{\mathsf{R}, \mathsf{G}, \mathsf{B}\}) \ , \tag{1}$$

where γ denotes non-linear scaling by the camera, and c denotes one of the color channels red (R), green (G)and blue (B). Typically, one assumes that the effect of γ can be reversed by raising the observed pixel to the power of $1/\gamma$. Then, the i_c is straightforward to neutralize (i. e., "change to white") multiplying the equation with $1/i_c$.

If only a single image is available, the recovery of i_c is an ill-posed problem. As each observed color consists of two unknowns, the material color and the light color. Hence, color constancy algorithms require additional constraints or assumptions to obtain a solution.

Figure 1: Example scenes under non-uniform illumination.

To make the problem better tractable, most of the previous work assumed homogeneous, globally uniform illumination. With this assumption, all pixels and color channels in the image can be used to estimate a single three-component vector of the illuminant color. This estimate is either based on image statistics, physics-based reasoning, or perceptual assumptions, see e. g., [BCF02, BMCF02, GGvdW11, TNI04, REA09].

Unfortunately, in real-world scenes, the assumption of a homogeneous uniform illumination is often greatly violated. Typical scenes often contain multiple light sources, like indoor and outdoor illumination, flash light and environmental lights. Figure 1 shows two such examples, where two light sources mix across the scene. Hence, the color and spatial distribution of each illuminant must be estimated, which makes the color constancy problem considerably harder.

So far, only few algorithms have been proposed to approach this non-uniform illumination. The earliest work is by Gijsenij *et al.* [GLG12] and Bleier *et al.* [BRB+11], which is a statistics-driven approach that has been demonstrated to operate well on synthetic and laboratory images. Other work showed good qualitative results on real images, but lacks a clear optimization criterion when estimating the spatial distribution of the light sources [REA11].

In this work, we extend [REA11] by investigating a physics-based multi-illuminant estimator. We obtain illuminant estimates for spatially limited, small locations in the image. Then, we filter these estimates to end up with two or three representative illuminant colors using a simple combinatorial objective function.

2 Illuminant Colors and their Spatial Distribution

The algorithm can be subdivided into three steps. First, the image is subdivided into smaller regions of approximately equal size. Then, the illuminant color is estimated locally on each region. In a third step, we seek a cover of some of these illuminant estimates across the scene that is optimal with respect to a given objective function.

Image sizes are normalized to a height of 480 pixels. Similar to [REA11], we first segment the image into similar colors using the algorithm by Felzenszwalb and Huttenlocher [FH04]. As some of these segments can still be quite large, we intersect this segmentation with a regular grid of 32×32 pixels. On each of these segments, we apply a physics-based illuminant-color estimator [REA11] to obtain a per-segment illuminant color estimate $s \in S$. The working principle of this estimator is to project the color values of the pixels into inverse-intensity chromaticity (IIC) space [TNI04]. Let $\mathbf{p} = (p_\mathsf{R} p_\mathsf{G} p_\mathsf{B})^\mathsf{T}$ an image pixel. The mapping to IIC space is for each pixel of the region of interest and each color channel $c \in \{\mathsf{R}, \mathsf{G}, \mathsf{B}\}$

$$p_c \rightarrow \left(\frac{1}{p_\mathsf{R} + p_\mathsf{G} + p_\mathsf{B}}, \frac{p_c}{p_\mathsf{R} + p_\mathsf{G} + p_\mathsf{B}} \right) \quad . \tag{2}$$

This projection is performed for each color channel c.

In IIC space, the color of the illuminant is obtained from analyzing pixel distributions after the projection according to Eqn. 2 that stem from the same underlying material, but expose different degrees of specularity. In real-world images, it is difficult to determine changes of material or specularness in pixels in an automated manner. To accomodate for that, it is reasonable to impose further constraints on the pixel distributions in IIC space. In prior work [REA11], it has been shown that heuristic shape constraints on the pixel distributions in IIC space can be used. Partially specular single-material pixels form a wedge shape in IIC space. The tip of the wedge points towards the y axis. Estimating the point of intersection of the wedge tip with the y axis yields an illuminant color estimate [TNI04, REA11]. We performe this estimation on the subdivided image, i. e. on image segments s, such that we eventually have several illuminant color estimates from different positions in the image.

The performance of the illuminant color estimator depends on the fact that at least some specular reflectance is present in the pixel distribution. To address this point, we compute for each segment s a simple confidence measure c_s. We estimated the degree of specular reflection per segment using an algorithm by Tan and Ikeuchi [TI05]. One of the assumptions of the specularity estimation is white illumination, which leads to a chicken-and-egg problem. However, we found the estimated specularity maps to be sufficiently good for this work. The confidence c_s corresponds to the sum of the per-pixel specularity weights, normalized over the image between 0 and 1.

Figures 2(a) and 2(b) show local estimation results. Black regions denote areas where the illuminant color estimator did not provide a reliable estimation result. The estimates still contain several apparent outliers. With the subsequent processing steps, we aim to remove the outliers, and to obtain a spatial distribution of 2 to 3 likely illuminants. Let $\mathbf{i}_s$ the illuminant estimate in segment s. We randomly select 200 instances of $\mathbf{i}_s$ from the local illuminant estimates. If less than 200 estimates are available, we use all. Let $\mathcal{Q}$ denote all subsets of these illuminant estimates containing 1, 2 or 3 illuminants, and $Q \in \mathcal{Q}$ one particular selection of illuminants. Then, we seek an optimal subset Q_opt subject to

$$Q_\mathsf{opt} = \underset{Q}{\mathsf{argmin}} \sum_{s \in S} \Big(1 - \sum_{\mathbf{q} \in Q} d(\mathbf{i}_s, \mathbf{q}) \Big)^2 c_s \quad , \tag{3}$$

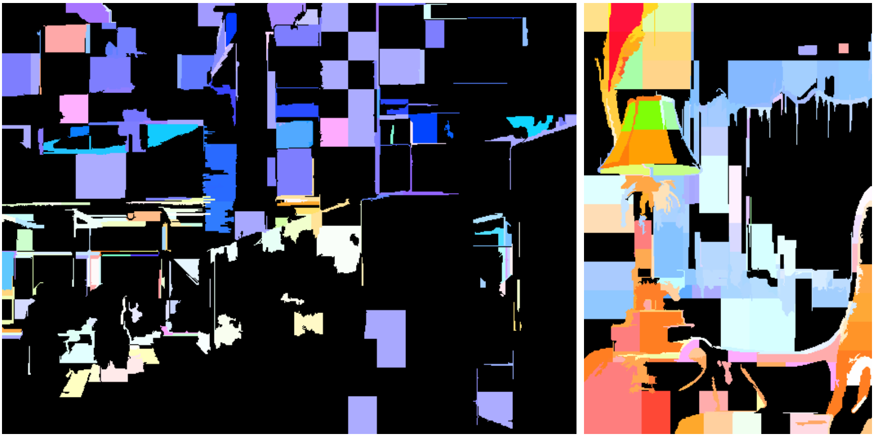

(a) (b)

Figure 2: Local illuminant estimates on the images in Fig. 1.

where $d(\mathbf{i}_s, \mathbf{q})$ is a piecewise linear error function, defined on the angular distance $a(\mathbf{i}_s, \mathbf{q}) = \cos^{-1}(\mathbf{i}_s \circ \mathbf{q})$ as

$$d(\mathbf{i}_s, \mathbf{q}) = \begin{cases} 1 - 9 \cdot a(\mathbf{i}_s, \mathbf{q}) & \text{if } a(\mathbf{i}_s, \mathbf{q}) \leq 0.1 \\ 0.1 - (a(\mathbf{i}_s, \mathbf{q}) - 0.1)/9 & \text{otherwise} \end{cases} . \tag{4}$$

Hence, we look for an 'optimal' set of illuminants, in the sense that most of the scene patches with high confidence high-confidence c_s are close to one of the candidate illuminants q.

Once q has been determined, we assign each illuminant color estimate $\mathbf{i}_s$ the color of the closest illuminant in S_{opt}. The outcome for the two example images is shown in Fig. 3. In comparison to Fig. 2(a) and Fig. 2(b), most of the outliers are suppressed, while still maintaining some scene structure in the illuminant estimates.

3 Results

Quantitative ground truth on multi-illuminant scenes is challenging to obtain, which is why our study is limited to qualitative results. We show and discuss two more examples in Fig. 4 and Fig. 5.

In Fig. 4, a slightly reddish spotlight illuminates the central part of the scene, while the background light is apparently blue. The painting in the background leads to many outlier illuminant estimates (see Fig. 4 top right). However, these outliers exhibit a wide variety of colors, and hence do not have much impact on the selection of the final illuminants. As Fig. 4 (bottom) shows, blue and a slightly yellowish tone are eventually selected. The spatial distribution of the filtered illuminant estimates is very reasonable, in the sense that it pretty much resembles the image structure.

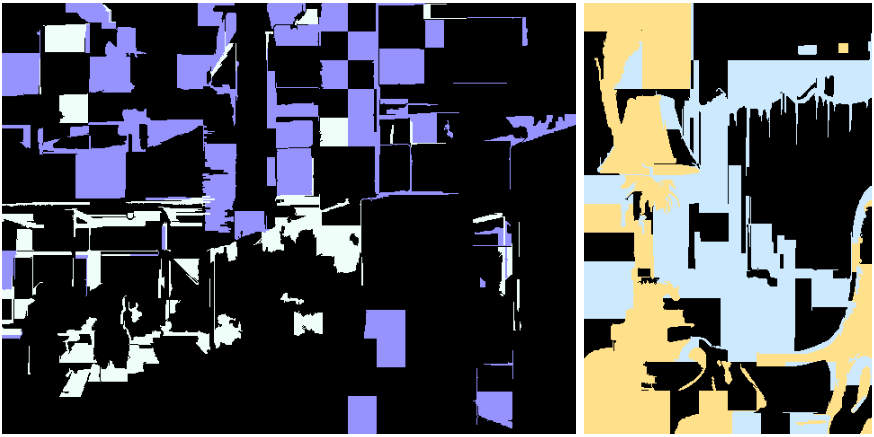

(a) (b)

Figure 3: Filtered illuminant estimates after computing Q_{opt}. Outliers are greatly suppressed.

This, however, is not always the case. Figure 5 contains a failure case. The subway scene is very cluttered, with several man-made illuminants and specular surfaces in the scene. The local illuminant estimates also mostly consist of outliers, scattered across the scene (see Fig. 5, top right). In this experiment, we restricted the number of estimated illuminants to exactly 2 or exactly 3, which is shown in Fig. 5 bottom left and right, respectively. In both cases, a bluish and very light red are selected. However, the assignments of illuminant colors to the image segments does not exhibit much of the underlying image structure. As a consequence, white balancing algorithm that uses these estimates as seed points might have difficulties to achieve a satisfying result.

4 Discussion

The presented algorithm demonstrates in several examples good qualitative performance. Initial experiments are encouraging, and suggest that it might be possible to use these estimates as seed points for a white balancing algorithm.

However, there are also some intrinsic limitations of the presented approach, that should be addressed in future work. First, the reasoning on the best illuminants is entirely done on the illuminant colors, i. e. the image content or spatial distribution is not taken into consideration.

Additionally, the optimization criterion implicitly favors solutions that cover a certain area in RGB-space. Assume a scenario where the true illuminants are very close to each other in RGB-space. Assume also that most of the illuminant estimates are distributed closely around the true illuminants, with some small added noise. In this case, the algorithm would most likely select illuminants that are further apart, just to cover also some of the more noisy estimates. In other words, it does not explicitly search for centers

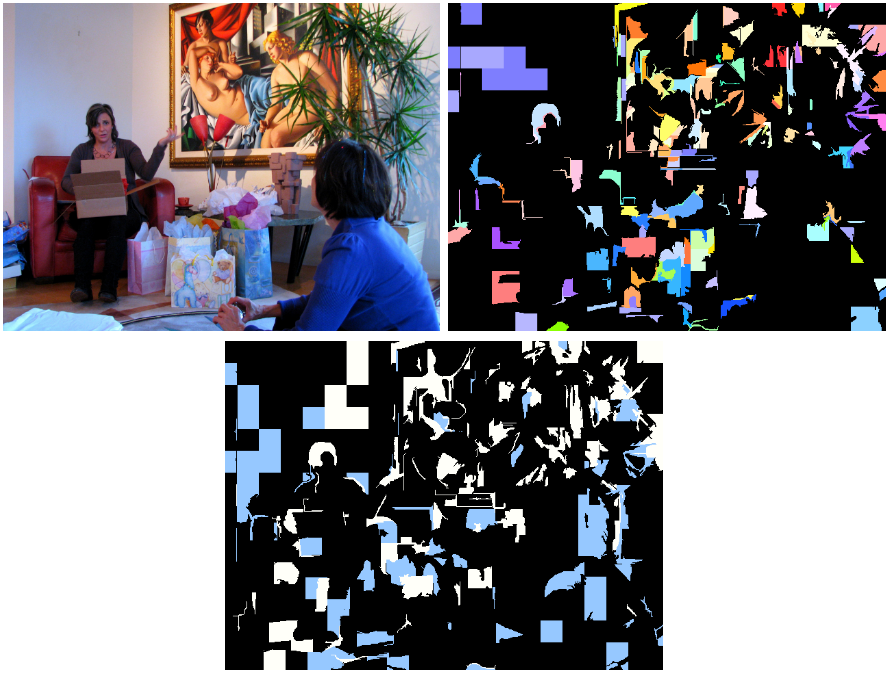

Figure 4: The painting in the background of the scene leads to many outlier illuminant estimates (top right). However, the relatively strict constraint on 2 or 3 illuminants discards illuminant estimates that are scattered across RGB space (bottom).

in densities, but blindly optimizes distances to bags of illuminant color estimates.

However, it is a relatively straightforward algorithm, that apparently operates effectively in many situations. In that sense, it has some attractivity in its own right due to its simplicity.

In future work, we will investigate ways to incorporate the image content for optimizing the color of the illuminants. We also plan to use such sparse estimates for seeding a white-balancing algorithm.

5 Conclusions

We presented a selection method for multiple local illuminant estimates. It chooses two to three reference illuminants in a multi-illuminant scene and recolors the estimates. The output of the algorithm can be used, e. g., for seeding an advanced white balancing algorithm. The algorithm itself is straightforward to implement, and given its simplicity, apparently performs reasonably well.

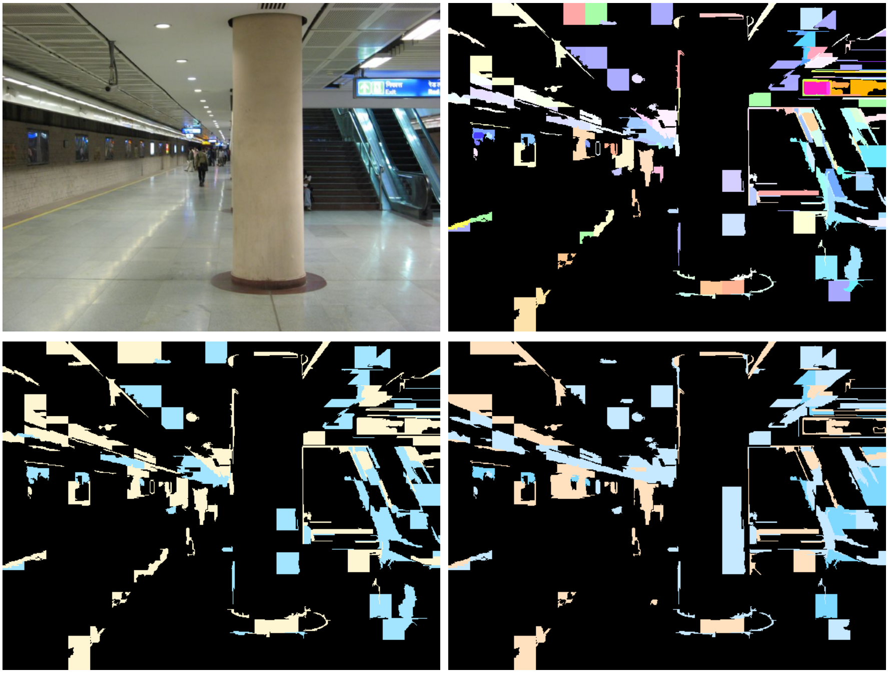

Figure 5: The subway station contains a quite complex lighting situation. Besides a lot of outliers, several details are also picked up by the local illuminant estimates (top right). The selection of 2 (bottom left) or 3 (bottom right) illuminants favors blue and white lights, but these illuminants are scattered across the scene, i. e., reveal not much scene structure.

References

[BCF02] K. Barnard, V. Cardei, and B. Funt. A Comparison of Computational Color Constancy Algorithms – Part I: Methodology and Experiments With Synthesized Data. *IEEE Transactions on Image Processing*, 11(9):972–983, September 2002.

[BMCF02] K. Barnard, L. Martin, A. Coath, and B. Funt. A Comparison of Computational Color Constancy Algorithms – Part II: Experiments With Image Data. *IEEE Transactions on Image Processing*, 11(9):985–996, September 2002.

[BRB+11] Michael Bleier, Christian Riess, Shida Beigpour, Eva Eibenberger, Elli Angelopoulou, Tobias Tröger, and André Kaup. Color Constancy and Non-Uniform Illumination: Can Existing Algorithms Work? In *IEEE Color and Photometry in Computer Vision Workshop (CPCV 2011)*, pages 774–781, November 2011.

[FH04] P. Felzenszwalb and D. Huttenlocher. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision*, 59, September 2004.

[GGvdW11] Arjan Gijsenij, Theo Gevers, and Joost van de Weijer. Computational Color Constancy: Survey and Experiments. *Transactions on Image Processing*, 20(9):2475–2489, September 2011.

[GLG12] Arjan Gijsenij, Rui Lu, and Theo Gevers. Color Constancy for Multiple Light Sources. *IEEE Transactions on Image Processing*, 21(2):697–707, February 2012.

[REA09] C. Riess, E. Eibenberger, and E. Angelopoulou. Illuminant Estimation by Voting. Technical report, Friedrich-Alexander University Erlangen-Nuremberg, September 2009.

[REA11] C. Riess, E. Eibenberger, and E. Angelopoulou. Illuminant Color Estimation for Real-World Mixed-Illuminant Scenes. In *Color and Photometry in Computer Vision Workshop*, pages 782–289, November 2011.

[TI05] R. Tan and K. Ikeuchi. Separating Reflection Components of Textured Surfaces Using a Single Image. *Transactions on Pattern Analysis and Machine Intelligence*, 27, February 2005.

[TNI04] R. Tan, K. Nishino, and K. Ikeuchi. Color Constancy Through Inverse-Intensity Chromaticity Space. *Journal of the Optical Society of America A*, 21, March 2004.