Illuminant Segmentation in Non-uniformly Lit Scenes

Cong Phuoc Huynh[1] Antonio Robles-Kelly[1,2]
[1]*National ICT Australia*,* *Locked Bag 8001, Canberra ACT 2601, Australia*
[2]*Research School of Engineering, Australian National University, Canberra ACT 0200, Australia*

Abstract

In this paper, we present a method for segmenting illuminants in non-uniformly lit scenes. Here, we view the illuminant colour at an image location as a mixture of the segmented illuminants. Based on the dichromatic structure of the image radiance space, we perform soft-clustering on the set of dichromatic planes corresponding to the neighbourhoods of pixel-sites in the image. We solve the soft-clustering problem with a deterministic annealing approach where the cost function is formulated based on the maximum entropy principle. We show results on real-world imagery and provide comparisons to an alternative method.

1. Introduction

The appearance of an object in a scene depends greatly on the illuminant power spectrum and the scene geometry. Hence, identifying the light sources in a scene can greatly benefit recognition and identification methods for machine vision applications based on photometric invariants [8, 6].

However, the recovery and identification of the illuminants in the scene has proven a difficult task in uncontrolled real-world imagery. This is mainly due to the fact that the recovery of the illuminants from a single image is an under-constrained problem [2]. To this end, Finlayson *et al.* [4] detected the illumination in the scene making use of a chromagenic camera. Wang and Samaras [11] detected and estimated the illuminants in the scene making use of a recursive least squares method. In [12], the authors detected the direction of the light sources by observing that the change in image intensities is maximum when the illuminant direction is perpendicular to the normal of the surface. Ebner [3] employed the local space average color to perform colour constancy irrespective of the illuminants used in the scene. Barnard *et al.* [1] have addressed the colour constancy problem making use of the information conveyed by the reflectance and the illuminants across non-uniformly lit scenes.

In this paper, we aim to label regions illuminated by distinct light sources as a preprocessing step for the recovery of spatially-varying illumination and photometric invariants in a scene. Here, we note that, when the scene is illuminated by multiple light sources, the illuminant colour at an image location can be viewed as a mixture of these sources. With this in mind, the illuminant segmentation problem amounts to finding their mixture coefficients per pixel. We exploit the dichromatic structure of the image radiance [9] to formulate the problem as a soft clustering one. Our formulation takes into account the probability of image locations being illuminated by an individual light source in a two step fashion. The first step consists of the extraction of a set of dichromatic planes as described in Section 2.1. Once these planes are obtained, we cluster them into groups, each of which intersects at a common vector representing an individual illuminant, as presented in Section 2.2.

2. Illuminant Segmentation

To commence, consider the image whose radiance, illumination and reflectance at the pixel u and the colour channel $k \in \{R, G, B\}$ are denoted as $I_k(u)$, $L_k(u)$ and $S_k(u)$, respectively. As mentioned earlier, to segment the illuminants in a scene lit by multiple light sources, we employ the dichromatic reflection model in [9]. The model can be expressed in a compact form as follows

$$\mathbf{I}(u) = g(u)\mathbf{L}(u) \bullet \mathbf{S}(u) + k(u)\mathbf{L}(u), \qquad (1)$$

where $\mathbf{I}(u)$, $\mathbf{L}(u)$ and $\mathbf{S}(u)$ are the radiance, illumination and reflectance vectors whose k^{th} entry corresponds to the colour channel $\mathcal{C}_k$, $g(u)$ and $k(u)$ are the shading factor and the specular coefficients and $\bullet$ denotes the element-wise vector product.

*NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

Note that the first term on the right-hand side of Equation 1 depends on both the material reflectance and the illuminant while the second term is only dependent on the illuminant. Therefore, the radiance vectors at surface locations of the same material and illuminated by the same illuminant belong to a dichromatic plane [5] spanned by the illuminant vector $\mathbf{L}(u)$ and the diffuse radiance vector $\mathbf{D}(u)$ of the material, where $\mathbf{D}(u) \triangleq \mathbf{L}(u) \bullet \mathbf{S}(u)$.

2.1. Extracting Dichromatic Planes

As a result, the radiance in each image can be viewed as a set of dichromatic planes. Moreover, those corresponding to the same light should form a cluster that intersects at the illuminant vector. Note that, if the scene illumination is piece-wise constant, *i.e.* each image pixel is illuminated by a single illuminant, the radiance at each pixel should lie exactly in one of these planes. However, in general, a scene location can be simultaneously illuminated by a mixture of light sources. Therefore, the pixel radiance does not necessarily lie on a single dichromatic plane.

Here, we can cast the problem of segmenting the illuminants in an image as a soft-clustering problem on the dichromatic planes containing the illuminant vectors. To this end, we extract a set of dichromatic planes from the image, each for the local neighbourhood at each pixel. This operation is undertaken with the assumption that the local neighbourhood is made of the same material and illuminated by the same light. This is not an unreasonable assumption since both the material and the illumination often vary smoothly across the image. In practice, these dichromatic planes can be obtained by performing Singular Value Decomposition (SVD) on the matrix whose columns are formed by the radiance vectors at the pixels in the neighbourhood under consideration [5, 10]. The basis vectors of the plane are, hence, given by the singular vectors corresponding to the two largest singular values of this matrix.

2.2. Clustering of the Dichromatic Planes

With the dichromatic planes for each pixel in the image at hand, we proceed to cluster them so as to segment the scene illuminants. Let us denote this set of dichromatic planes as $\mathcal{P}$. If a pixel is illuminated by an illuminant L, then L should lie on the dichromatic plane $P \in \mathcal{P}$ corresponding to that pixel. This observation justifies the use of the distance metric $d(L, P)$ between a dichromatic plane and the illuminant for the purpose of clustering. A zero-distance implies that the group of pixels belonging to the plane are purely illuminated by a single illuminant. The greater the distance, the lower the power of the illuminant L impinging upon the pixel under consideration.

We model the association between a dichromatic plane P and an illuminant L using a probability measure $p(L|P)$. Thus, we can view the illuminant clustering problem as finding a set of illuminant spectra $\mathfrak{L} = \{L\}$ and a set of association probablities $\mathcal{D} = \{p(L|P)|L \in \mathfrak{L}, P \in \mathcal{P}\}$ which minimise the total expected plane-illuminant distance

$$C = \sum_{P\in\mathcal{P}} \sum_{L\in\mathfrak{L}} p(L|P)d(L, P) \tag{2}$$

subject to the constraints $\sum_{L\in\mathfrak{L}} p(L|P) = 1$ and $\|L\| = 1$.

In Equation 2, we view $d(L, P)$ as the orthogonal distance between the vector L and the plane P. Suppose that the two known basis vectors spanning the dichromatic plane P are $\mathbf{z}_1(P)$ and $\mathbf{z}_2(P)$. The linear projection matrix $Q(P)$ onto P is computed using the relation $Q(P) = A(P)(A(P)^T A(P))^{-1} A(P)^T$, where $A(P) = [\mathbf{z}_1(P), \mathbf{z}_2(P)]$. The distance $d(L, P)$ is therefore formulated as $d(L, P) = \|L - Q(P)L\|^2$.

Note that the trivial solution to the cost function in Equation 2 is a nearest-neighbour one which associates dichromatic planes to their closest illuminant spectrum vectors with a unit probability. To formulate the problem as a soft-clustering one, we enforce an additional constraint on the distribution of association probabilities $\mathcal{D} = \{p(L|P)|L \in \mathfrak{L}, P \in \mathcal{P}\}$ based on the maximum entropy principle [7].

The uncertainty of the plane-illuminant association probability is given by

$$H(\mathcal{D}) = -\sum_{P\in\mathcal{P}} \sum_{L\in\mathfrak{L}} p(L|P) \log p(L|P) \tag{3}$$

Making use of the entropy, the cost function becomes $C_{Entropy} = C - \mathcal{L}$ where

$$\begin{aligned} \mathcal{L} = TH(\mathcal{D}) &+ \sum_{P\in\mathcal{P}} \alpha(P) \left(\sum_{P\in\mathcal{P}} p(L|P) - 1 \right) \\ &+ \sum_{L\in\mathfrak{L}} \beta(L) \left(\|L\|^2 - 1 \right) \end{aligned} \tag{4}$$

In the equations above, $\alpha(P)$, $\beta(L)$ and $T \geq 0$ are Lagrange multipliers. While T weighs the level of randomness of the plane-illuminant association probabilities, the second term in Equation 4 enforces the total probability constraint for every dichromatic plane P and the third term imposes the norm constraint on L.

2.2.1 Deterministic Annealing Optimisation

Here, we turn our attention to the optimisation of the cost function presented above. To do this, we observe that the Lagrangian multiplier T can be viewed as the system temperature of an annealing process. Using this analogy, we employ a deterministic annealing approach to minimising the cost function by varying the temperature from an initial high value to a zero value at the end of the process. The process converges to a thermal equilibrium whereby the system undergoes a "phase transition" as the temperature is lowered and the optimal solution is "tracked" through to the cooling period. At zero temperature, we can directly minimise the total expected plane-illuminant distance so as to obtain the final plane-illuminant association probabilities and the illuminant clusters.

Note that the annealing process described above is similar to a soft-clustering one where, at the initial high temperature, all the dichromatic planes are assumed to be spanned by a single illuminant. As the temperature drops, the set of illuminants grows in size. This process passes through several "phase transitions", at which new illuminants are split from the existing ones. At each temperature T, the algorithm alternates between two interleaved minimisation steps until it converges. These steps aim at finding the optimal plane-illuminant association probabilities for the current set of illiminants. The optimisation of these variables is described in the following sections.

2.2.2 Estimating the Illuminant Association Probability

At a constant temperature , we fix the set of illuminants and estimate the plane-illuminant association probabilities that minimise the cost function $C_{Entropy}$. The minimisation is effected by setting the partial derivative of $C_{Entropy}$ with respect to $p(L|P)$ to zero, which yields

$$p(L|P) = \exp\left(\frac{-d(L,P)}{T} + \frac{\alpha(P)}{T} - 1\right) \forall L, P \quad (5)$$

Since $\sum_{L\in\mathfrak{L}} p(L|P) = 1$, the optimal association probability distribution for a fixed set of illuminants $\mathfrak{L}$ is given by the Gibbs distribution

$$p(L|P) = \frac{\exp\left(\frac{-d(L,P)}{T}\right)}{\sum_{L'\in\mathfrak{L}} \exp\left(\frac{-d(L',P)}{T}\right)} \quad (6)$$

2.2.3 Estimating the Illuminants

We now fix the plane-illuminant association probabilities and seek the optimal set of illuminants for the dichromatic planes extracted from the image. We commence by rewriting the distance metric $d(L,P)$ as $d(L,P) = \|L - Q(P)L\|^2 = \|(J - Q(P))L\|^2 = L^T R(P)^T R(P) L$, where J is the identity matrix and the matrix $R(P) = J - Q(P)$ is known. Since $R(P)^T R(P) = R(P)$, we can simply write $d(L,P) = L^T R(P) L$. Using this compact form, we can derive that $\frac{\partial d(L,P)}{\partial L} = 2R(P)L$. Substituting this relation into the derivative of $C_{Entropy}$ with respect to the illuminant L, we have

$$\frac{\partial C_{Entropy}}{\partial L} = 2\left(\sum_{P\in\mathcal{P}} p(L|P)R(P)L - \beta(L)L\right)$$

As a result, the optimal set of illuminants is found by setting $\frac{\partial C_{Entropy}}{\partial L}$ to zero. This yields

$$\sum_{P\in\mathcal{P}} p(L|P)R(P)L = \beta(L)L \quad (7)$$

From Equation 7, it is straightforward to deduce that L is an eigenvector of the matrix $M = \sum_{P\in\mathcal{P}} p(L|P)R(P)$.

3. Experiments

In this section, we illustrate the utility of our method for illuminant segmentation as compared to the ground truth and those yielded by Ebner's method [3]. Although the method in [3] mainly performs white balancing on images lit by spatially varying illumination, we can utilise the ratio of the original image to the white-balanced one as an indication of the spatial illumination variation. Here we adopted the variant of Ebner's method which operates in the original RGB space.

We performed our experiments on the data set of multi-illuminant images acquired by Bleier *et al.* [2]. For our experiments, we use the correctly exposed images of three scenes captured under two Reuter lamps with or without colour filters. The lamp positioned on the left side of the scene was mounted with one filter of a set of three LEE 201, 202 and 281 filters. Likewise, the one on the right side was mounted with one of three LEE 204, 205 and 285 filters. To acquire the ground truth illuminant for these images, the same scenes were spray-painted in gray and re-captured under the same illumination condition.

In Figure 1, we present the illuminant segmentation results. The two left-hand columns show the input images and the colour-coded ground truth illuminant regions with their foreground mask. In the right-hand columns, we show the illuminant segments in the corresponding regions detected by our algorithm and Ebner's method [3], where we have distinguished the segments from each other using high-contrast colours.

We observe that our method delivers segments in close accordance with the ground truth. Furthermore,

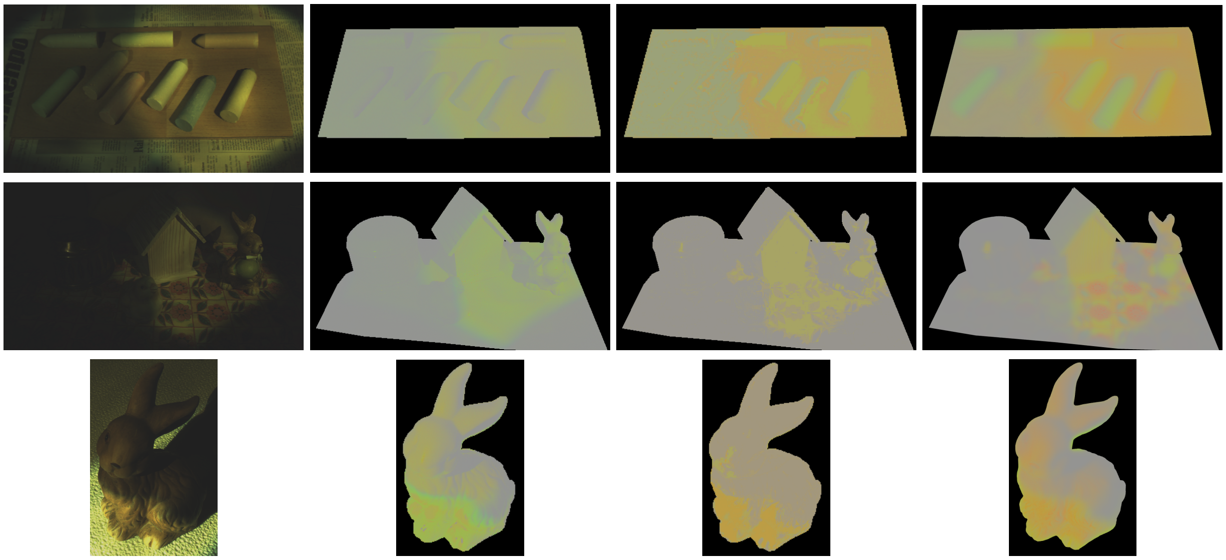

Figure 1. From left-to-right: Input images, the ground-truth illuminant maps, the illuminants segmented by our algorithm and those recovered using Ebner's method [3].

Ebner's method is more likely to misinterpret material segments as illuminant regions. Examples of this phenomenon is the green chalk regions under the left-hand side illuminant in the first scene and the printed flower regions in the second scene. This phenomenon is due to the nature of the alternative method. This is because it performs colour correction locally and thus its output is affected by the material reflectance. On the other hand, our method considers dichromatic planes across the whole image, which alleviates this problem.

4. Conclusions

We have presented a method for the segmentation of illuminants in scenes illuminated by multiple lights. The method hinges on the soft-clustering of the dichromatic planes recovered for every pixel in the image. We have formulated the soft-clustering problem with a constraint based on the maximum entropy principle. Subsequently, we have minimised the cost function with a deterministic annealing approach. Finally, we have performed experiments on real-world imagery and compared our results to those yielded by an alternative method.

References

[1] K. Barnard, G. D. Finlayson, and B. V. Funt. Color constancy for scenes with varying illumination. *Computer Vision and Image Understanding*, 65(2):311–321, 1997.

[2] M. Bleier, C. Riess, S. Beigpour, E. Eibenberger, E. Angelopoulou, T. Tröger, and A. Kaup. Color constancy and non-uniform illumination: Can existing algorithms work? In *ICCV Workshops*, pages 774–781, 2011.

[3] M. Ebner. Color constancy using local color shifts. In *European Conference on Computer Vision*, pages 276–287, 2004.

[4] G. D. Finlayson, C. Fredembach, and M. S. Drew. Detecting illumination in images. In *IEEE International Conference on Computer Vision*, 2007.

[5] G. D. Finlayson and G. Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *International Journal of Computer Vision*, 42(3):127–144, 2001.

[6] D. W. Jacobs, P. N. Belheumer, and R. Basri. Comparing images under variable illumination. In *Computer Vision and Pattern Recognition*, pages 610–617, 1998.

[7] E. T. Jaynes. Information Theory and Statistical Mechanics. *Phys. Rev.*, 106(4):620–630, 1957.

[8] S. Nayar and R. Bolle. Reflectance based object recognition. *International Journal of Computer Vision*, 17(3):219–240, 1996.

[9] S. A. Shafer. Using color to separate reflection components. *Color Research and Applications*, 10(4):210–218, 1985.

[10] S. Tominanga and B. A. Wandell. Standard surface-reflectance model and illuminant estimation. *Journal of the Optical Society of America A*, (6):576–584, 1989.

[11] Y. W. and D. Samaras. Estimation of multiple illuminants from a single image of arbitrary known geometry. In *European Conference on Computer Vision*, pages 272–288, 2002.

[12] Y. Zhang and Y.-H. Yang. Multiple illuminant direction detection with application to image synthesis. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(8):915–920, 2001.