

Colour constancy based on Bayesian inference on scene statistics

S. Nakauchi, T. Toyota and H. Honjyo

*Department of Information & Computer Sciences, Toyohashi University of Technology
Toyohashi 441-8580 (JAPAN)*

Corresponding author: S.Nakauchi (naka@bpel.ics.tut.ac.jp)

ABSTRACT

The problem of colour constancy is formulated as Bayesian inference on scene statistics to recover the colour of an incident illumination. We begin with the data rendering equation describing the relation between the illuminant and statistics of surface colours, and statistics of the observed scene. We here focus on the first and second-order statistics (mean and luminance-colour correlation) of the scene which are likely to be cues for the illuminant colour as suggested by recent psychophysical observations. Then, we construct prior distributions for statistics of the illuminants and a set of surface colours describing particular illuminants and surfaces existing in the world. Simulation results show that Bayesian estimates for the illuminant colour which maximize a posteriori probability computed with a given scene are robust across hue distribution changes of surface colours.

1. INTRODUCTION

From a computational point of view, the problem of colour constancy is conventionally defined as recovering illuminant colour and/or surface colours from a set of sensory responses. This is one of well-known ill-posed problems in vision because both illuminant and surface colours are generally unknown and there exist infinite combinations giving the same retinal image to be observed. Many researches have tried to deal with the colour constancy problem by making additional assumptions about the world. Among them, the most traditional and widely-used one is the so-called gray-world algorithm¹ assuming that the spatial average of surface colours in a scene is achromatic. In other words, mean colour of a given scene is assumed to give the colour of the incident illumination. Although this assumption is simple and can be closely related to colour adaptation mechanism, ‘gray-world assumption’ is often violated in real scenes, and thus, it does not work well. Even in such a case, as an extensive literature pointed out, human vision exhibits (approximate) colour constancy suggesting that human vision may use different assumptions other than or in addition to the gray-world assumption.

Recent works^{2,3} have demonstrated that the second-order statistics of a scene systematically modulate colour appearance. Regarding the colour constancy problem, Golz-MacLeod² showed that gray settings systematically depends on the luminance-colour (redness in their case) correlation and they interpreted this result as indicating that the visual system uses the second-order statistics as a cue for the illuminant colour estimation. These findings do not necessarily mean that the gray-world as the first-order statistics does not help to achieve colour constancy, but the visual system may infer the illuminant or surface colours by combining these multiple cues in a given scene.

This article provides a Bayesian approach to treat the colour constancy problem as probabilistic inference on the illuminant colour by combining cues from the first and second-order statistics of a given scene. Simulation results show that the Bayesian estimates for the illuminant colour are robust despite the hue distributions of surface colours suggesting that the luminance-colour correlation plays an important role.

2. METHOD

The first step of the Bayesian approach to the colour constancy problem is describing the data rendering equation. In a conventional physics-based approach⁴, this may correspond to the image formation which describes how the spectral surface reflectance and spectral power distribution of illuminant relate to the sensor response at each location. Our approach, on the other hand, begins with description of the relation among the illuminant colour, statistics of a set of surfaces, and statistics of

the observed scene as the data rendering equation $\mathbf{y} = f(\mathbf{x})$, where $\mathbf{x} = (\mathbf{x}^E, \mathbf{x}^S)^t$ is a vector representing the illuminant colour and the statistics of a set of surfaces and $\mathbf{y} = (\bar{r}, \bar{b}, C_{rY}, C_{bY})^t$ is a vector representing the first and second-order statistics of the observed scene. $\mathbf{x}^E = (r_E, b_E)^t$ is the chromaticity coordinate of the illuminant, $\mathbf{x}^S = (\bar{r}_S, \bar{b}_S, C_{rY}^S, C_{bY}^S)^t$ is the first and second-order statistics (mean and luminance-colour correlation) of a set of surfaces under an equal-energy white illuminant providing the information about the surface reflectance. These statistics are calculated for the chromaticity coordinates defined in a MacLeod-Boynton diagram⁵ (in a logarithmic scale).

A multiple regression model was used to construct the rendering equation $\mathbf{y} = f(\mathbf{x})$. Model parameters were estimated for a set of \mathbf{x}, \mathbf{y} data obtained from various simulated scenes. In scene simulations, we used spectral reflectance data set of Munsell colour chips⁶ and artificially generated illuminants of 94 different colours represented by sinusoidal-shaped power distribution with variable amplitude and phase. Each scene consists of 225 surfaces randomly selected from all hue to generate ‘gray-world’ (GW) scenes and from a particular 10 different hue distributions (R, Y, G, B, P, YR, PB, GY, BG and RP) to generate ‘coloured’ scenes. For each hue distribution, 20 scenes were randomly generated, thus, 20680 scenes in total (94 illuminant colours \times 11 hues \times 20).

Figure 1 illustrates the training data set for the model: (a) chromaticity of illuminant $(r_E, b_E)^t$, (b) mean chromaticity $(\bar{r}_S, \bar{b}_S)^t$ and (c) luminance-colour correlation $(C_{rY}^S, C_{bY}^S)^t$ of a set of surfaces, (d) mean chromaticity $(\bar{r}, \bar{b})^t$ and (e) luminance-colour correlation $(C_{rY}, C_{bY})^t$ of the observed scenes. For a display reason, we showed only one scene for each hue distribution. Mean surface colour and illuminant colour are represented by different symbols and colours in each panel, respectively. The problem of colour constancy is formulated here as recovering (a), (b) and (c) from (d) and (e). It should be noted that mean chromaticity of the observed scene varies depending on both the mean surface colours and illuminant colours (Fig.1(d)), while the luminance-colour correlation of the observed scene is more robust to the illuminant colour changes (Fig.1(e)). This property is supposed to play an important role in the proposed model, that is, the luminance-colour correlation can provide reliable information about statistics of the surface colour set despite the illuminant changes.

Within the Bayesian framework, the data rendering equation corresponds to the likelihood as a probability distribution. We describe the likelihood $p(\mathbf{y} | \mathbf{x})$ as a normal distribution, that is, $p(\mathbf{y} | \mathbf{x}) = N(\mathbf{y}; f(\mathbf{x}), \Sigma)$ where $N(\cdot)$ is a normal distribution with a mean of $f(\mathbf{x})$ and a covariance matrix of Σ which is estimated from a residual error of the model: $f(\mathbf{x}) - \mathbf{y}$. Note that the model could fit $\mathbf{y} = f(\mathbf{x})$ extremely well: correlation coefficients between \mathbf{y} and $f(\mathbf{x})$ were higher than 0.99.

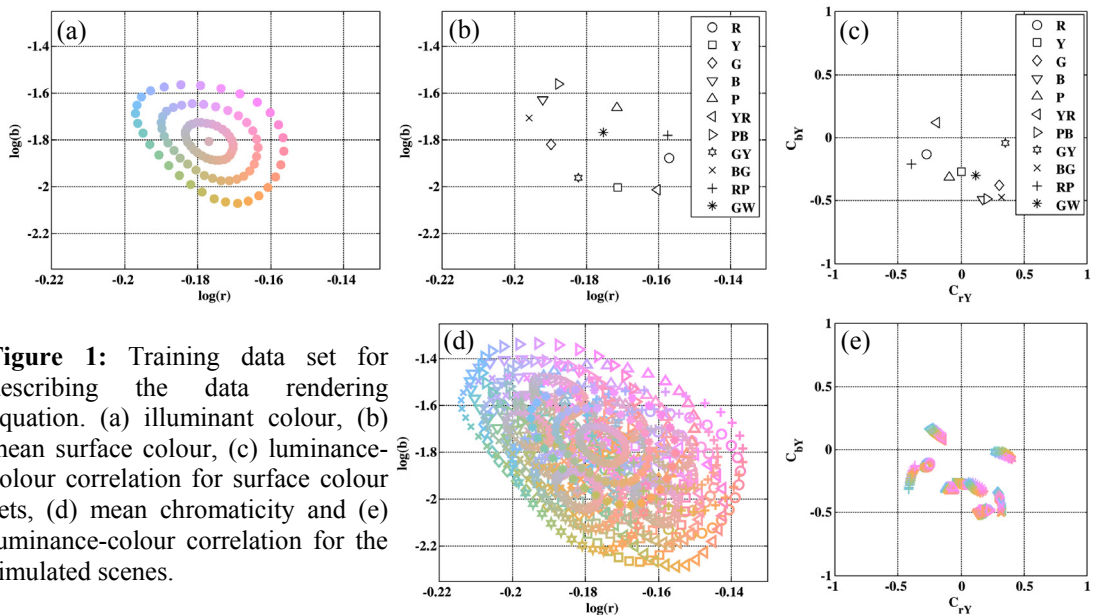


Figure 1: Training data set for describing the data rendering equation. (a) illuminant colour, (b) mean surface colour, (c) luminance-colour correlation for surface colour sets, (d) mean chromaticity and (e) luminance-colour correlation for the simulated scenes.

Finally, we describe the prior probability representing what is known about the illuminant and a set of surfaces before observing the scene. The prior about the illuminant colour $p(\mathbf{x}^E)$ was described by a two-dimensional normal distribution with the mean of white (achromatic) chromaticity and a certain variance fitted to the data used in the simulation shown in Fig. 1(a). The prior about the surface $p(\mathbf{x}^S)$ was described by a four-dimensional normal distribution which is also fitted to the data used in the simulation shown in Fig. 1(b) and (c). It can be assumed that the illuminant colour and a set of surfaces are independent. therefore, the total prior $p(\mathbf{x})$ was described as $p(\mathbf{x}) = p(\mathbf{x}^E)p(\mathbf{x}^S)$.

Given the likelihood $p(\mathbf{y} | \mathbf{x})$ and the prior $p(\mathbf{x})$, we can compute the posterior probability which means what is known about the illuminant and a set of surfaces after observing the scene. Using Bayes formula, the posterior can be $p(\mathbf{x} | \mathbf{y}) = p(\mathbf{y} | \mathbf{x})p(\mathbf{x}) / p(\mathbf{y}) \propto p(\mathbf{y} | \mathbf{x})p(\mathbf{x})$. Then, a MAP (maximum posterior) estimate $\tilde{\mathbf{x}}$ can be obtained by minimizing the expected loss function $\tilde{L}(\tilde{\mathbf{x}} | \mathbf{y}) = \int_{\mathbf{x}} -\delta(\tilde{\mathbf{x}} - \mathbf{x})p(\mathbf{x} | \mathbf{y})d\mathbf{x}$ for the observed scene statistics \mathbf{y} .

3. RESULTS

To evaluate the model performance, we first tested it for scenes similar to stimuli used by Golz-MacLeod². Figure 2 shows the estimated illuminant colours for scenes of which mean chromaticity is gray ($\log r, \log b$) = (-0.1769, -1.8069) and the luminance-red and luminance-blue correlation were varied from -0.9 to +0.9 at 0.1 intervals. As clearly seen, the estimated luminance colour systematically changes depending on the luminance-colour correlation. This simulation result resembles the observation by Golz-MacLeod, that is, more reddish illuminant is estimated for higher luminance-redness correlation C_{rY} . Furthermore, our model predicts that luminance-blueness correlation C_{bY} also affects the estimated illuminant colour in a similar fashion.

Next, we test how the luminance-colour correlation of the scene helps to achieve the colour constancy. We compared the performance for illuminant colour estimation among the gray-world algorithm, the proposed Bayesian model described above, and the Bayesian model without luminance-colour correlation term, that is, \mathbf{x} and \mathbf{y} are defined as $\mathbf{x} = (r^E, b^E, \bar{r}_S, \bar{b}_S)^T$ and $\mathbf{y} = (\bar{r}, \bar{b})^T$. Scenes to be tested were generated using sets of surfaces with 11 different hue distributions as in the training data set and illuminant of 5 different colours (red, green, blue, yellow and white). 20 scenes were generated for each combination of hue distribution and illuminant colour, thus 1100 scenes in total (11 hues \times 5 illuminant colour \times 20).

Figure 3 displays the chromaticity coordinates of the estimated illuminant colours by (a) gray-world algorithm, (b) Bayesian without correlation, and (c) proposed method. Shape and colour of symbols represent the hue distribution of surfaces and the illuminant colour, respectively. Filled circles are true illuminant colours.




Figure 2: Estimated illuminant colour for various luminance-colour correlations with the fixed mean chromaticity. Result resembles the observations by Golz-MacLeod².




Figure 3: Estimated illuminant colour by (a) Gray-world algorithm, (b) Bayesian without correlation, and (c) proposed method. Shape and colour of symbol represent the hue distribution of surfaces and the illuminant colour, respectively. Filled circles are true illuminant colours.

gather at the filled circles which represent ‘true’ illuminant colours. Although no models achieved perfect colour constancy, performance of the proposed model is superior to other two models because estimates tend to be clustered according to illuminant colour despite the hue distributions of surfaces.

To evaluate the performance quantitatively, we analyzed the variance of estimates for illuminant colour. Total variance of estimates is divided into two components, the variance due to the illuminant colours (V_E) and one due to the hue distributions of surfaces (V_S). We here focus on the ratio of these two components. According to the definition of colour constancy, ratio of variance V_E/V_S is expected to show larger values for a better model. As shown in Figure 4, the proposed model is superior to other two models, that is, sensitivity of the estimates for the illuminant colour were about 2.5 times as high as that for the hue distribution of surfaces. On the other hand, the gray-world and the Bayesian without correlation show similar performance. This is because both models use only the first-order statistics (mean) although the constraint on surface colours is slightly different (hard constraint in gray-world and soft constraint in Bayesian). This result shows that the luminance-colour correlation can be a crucial cue for achieving colour constancy.




Figure 4: Variance ratio V_E/V_S of estimates for illuminant colour by gray-world, Bayesian without correlation and proposed model.

4. CONCLUSIONS

This article demonstrated that the colour constancy problem can be described as probabilistic inference on the scene statistics within the Bayesian framework, and indicated that the luminance-colour correlation, in addition to the mean colour of the scene, may play an important role for colour constancy. Concept of the Bayesian inference suits the idea⁷ that the visual system incorporates implicit knowledge of the environment and image formation to infer object properties from the ambiguous images. Colour constancy is such a typical vision problem to recover the unknown illuminant colour (or surface colours) by incorporating the implicit knowledge of the illuminant colour and statistics of surface colors. Quantitative comparison with other related works based on the correlation⁸ and the experimental investigations⁹ on how the luminance-colour correlation affect the illuminant colour estimation are remained as future works.

References

1. G. A. Buchsbaum, "Spatial processor model for object colour perception," *J. Franklin Inst.* 310, 1-26 (1980).
2. J. Golz and D. I. MacLeod, "Influence of scene statistics on color constancy," *Nature* 415, 637-645 (2002).
3. R. Mausfeld and J. Andres, "Second-order statistics of colour codes modulate transformations that effectuate varying degree of scene invariance and illumination invariance," *Perception* 31, 299-224 (2002).
4. D. H. Brainard and W. T. Freeman, "Bayesian color constancy," *J. Opt. Soc. Am. A* 14, 1393-1411 (1997).
5. D. I. MacLeod and R. M. Boynton, "Chromaticity diagram showing cone excitation by stimuli of equal luminance," *J. Opt. Soc. Am. A* 69, 1183-1186 (1979).
6. Joensuu Spectral Database, <http://spectral.joensuu.fi>.
7. W. S. Geisler and D. Kersten, "Illusions, perception and Byes," *Nature Neurosci.* 5, 508-510 (2002).
8. G. D. Finlayson, S. D. Hordley and P. M. Hubel, "Color by correlation: A simple, unifying framework for color constancy," *IEEE Trans. Pattern Anal. Mach. Intell. PAMI-23*, 1209-1221 (2001).
9. J. J. M. Granzier, E. Brenner, F. W. Cornelissen and J. B. J. Smeets, "Luminance-color correlation is not used to estimate the color of the illumination," *J. Vis.* 5, 20-27 (2005).