

Color constancy in natural scenes independent of an explicit illuminant cue

Kinjiro Amano, David H. Foster; Sensing, Imaging, and Signal Processing Group, School of Electrical and Electronic Engineering, University of Manchester; Manchester, UK

Sérgio M. C. Nascimento; Department of Physics, Gualtar Campus, University of Minho; Braga, Portugal

Abstract

The ability of observers to make reliable judgments of surface color despite changes in an illuminant (“color constancy”) has sometimes been attributed to their estimating the spectral properties of the illuminant in order to discount its effects. To test how information about the illuminant contributes to surface-color judgments, observers’ performance in two surface-color-matching experiments with natural scenes, acquired with a high-resolution hyperspectral imaging system, was subjected to a multivariate analysis of variance. In the first experiment, the sky was directly visible to the observer, and its color was varied. In the second experiment, a large gray sphere was introduced into the scene so that its illumination by the sun and sky was directly visible to the observer and the color of that illumination was varied. Although observers’ surface-color matches varied across conditions, there was no reliable effect of the illuminant cue. Even when the sky was invisible, performance did not worsen. Judging surface color in natural scenes seems not to require knowledge of the illuminant.

Introduction

Human observers can make reliable judgments of surface color despite changes in the spectral properties of an illuminant that is not in the field of view. Some theories of this color constancy assume that the observer estimates the spectral properties of the illuminant, in order to compensate for its effects, an estimation that is theoretically possible in some circumstances [1-4]. For the task of matching colored papers in Mondrian-like patterns, it appears to be unnecessary [5], but it is possible that in scenes where explicit knowledge about an illuminant is available rather than indirect inference, there might be an effect on observers’ judgments.

To test this hypothesis, two surface-color-matching experiments were performed with images of natural scenes obtained from high-resolution hyperspectral images [6]. The test surface in each scene formed part of the scene and an explicit cue to the illuminant was manipulated in different ways to bias the information about its spectrum. By contrast with analyses of observers’ performance summarized by a univariate constancy index [7], which may mask the effects of some stimulus manipulations, here CIE u' , v' values of observers’ matches were subjected to a multivariate analysis of variance (MANOVA).

In the first experiment, the sky (but not the sun) in the scene was clearly visible to the observers, and its color was varied. In the second experiment, a large gray sphere was introduced into a scene so that the illumination upon it (avoid preempting details in Stimuli) was also clearly visible, and the color of that illumination was varied. As detailed later, the

MANOVA revealed no significant effect of illuminant color in either scene, nor an effect of initial test surface color. Judging surface color in natural scenes seems to be independent of knowledge of illuminant color.

Methods

Apparatus

Stimuli were presented on the screen of a 21-inch RGB color monitor (Trinitron Color Graphic Display, model GDM-F500R, Sony Corp., Tokyo, Japan), with spatial resolution 1600×1200 pixels, and 10-bit intensity resolution on each RGB gun, controlled by a graphics workstation (Fuel V12, Silicon Graphics, Inc., CA, USA). A telespectroradiometer (SpectraColorimeter, PR-650, Photo Research Inc., Chatsworth, CA, USA), calibrated by the National Physical Laboratory, was used regularly to calibrate the display system. Errors in the displayed CIE (x, y, Y) coordinates of a white test patch were < 0.005 in (x, y) and $< 5\%$ in Y ($< 10\%$ at low light levels).

Stimuli

Two natural scenes (an urban scene in Experiment 1 and a rural scene in Experiment 2) were selected from a hyperspectral image database [6]. The scenes are illustrated in Fig. 1, but note that the highlights were heavy clipped. The urban scene consisted of a group of 16th–18th-century buildings and sky. The test surface was a roof in the middle of the image, which was illuminated by the sky. The rural scene consisted of fields and trees, some low farm buildings, distant mountains, and sky. The 30-cm-diam. gray sphere was introduced into the scene so that it reflected light directly from the sun. The test surface was the roof of one of the farm buildings in the lower centre of the scene. In both scenes the sky occupied about a quarter of the image, and was clearly visible to observers. The images on the screen of the monitor subtended approx. $18^\circ \times 13^\circ$ visual angle at a viewing distance of 100 cm. The test surfaces subtended approx. $1^\circ \times 0.5^\circ$ and $3^\circ \times 1^\circ$ for the urban and rural scenes, respectively. The gray sphere subtended approx. 1° .

Procedure

In each trial, two images of a particular scene were presented in sequence on the screen of the color monitor, each for 1 s, with no interval. The images differed in the global illuminants: for the urban scene of Experiment 1, the illuminant was first daylight with correlated color temperature 25000 K and then 6700 K; for the rural scene of Experiment 2, it was first 15000 K and then 5700 K because of the different color gamut of the scene and the limitations of the color monitor.



Figure 1. Sample images of natural scenes. Urban (a, b) and rural (c, d) images were used in Experiments 1 and 2, respectively. No manipulation of sky color and sphere illuminant (a, c), and “more red” manipulation (b, d). The test surface in each scene is indicated by and arrow.

The spectral reflectance of the test surface in the second image was changed randomly in a way quantified by an equivalent local change in daylight [8]; that is, the second global daylight was replaced by a spatially uniform local illuminant constructed from a linear combination of the daylight spectral basis functions [9]. The chromaticity of this local illuminant was sampled randomly in each trial from a convex gamut in the CIE 1976 (u' , v') diagram comprising 65 locations, including both the first and second global illuminants. Varying the chromaticity of this local illuminant is closely related to varying the chromaticity of the test surface. Technical details of the surface-reflectance manipulations are given elsewhere [10].

Observers reported in each experimental trial whether there was an illuminant change or an illuminant change accompanied by a change in the surface-color of the test surface. Responses were made with mouse buttons.

Changes to the spectra of the sky in the urban scene of Experiment 1 and of the light from the sun and sky on the gray reflecting sphere in the rural scene of Experiment 2 were made in the same way as the changes to the spectra of the test surface. In Experiment 1, four color shifts of the sky were selected, “neutral” (no change in color), “blue”, “red” and “more red”, and in Experiment 2, three color shifts of the light on the sphere were selected, “neutral”, “red” and “more red”. The “red” was located midway between the two global illuminants on the daylight locus in the CIE 1976 (u' , v') diagram, the “blue” was an equal and opposite distance beyond the first illuminant, and the “more red” coincided with the second global illuminant. Examples of the “more red” manipulation are shown in Fig. 1 (b) and (d). These visible illuminant changes were the same for the first and second

images, for making inconsistent changes to the two images is known to worsen color-constancy judgments. Independent of these variations, the initial color of the test surface was also varied so that it was initially “more yellow”, “more green”, etc. These variations were achieved by applying local scene illuminants whose colors differed in 0.015 steps along the u' , v' -axes.

In each experimental session, there was just one color change to the sky or sphere and just one test-surface color. In all, 20 conditions were tested for Experiment 1, and 15 for Experiment 2. In Experiment 2, a further control condition was introduced in which the scene was cropped to remove both the sphere and sky.

Images were presented in a dark surround, and the luminance at each pixel varied from 0 to 33 cd m^{-2} .

Observers

Twelve observers, aged 23–34 years, took part in the experiments: 2 male and 4 female for Experiment 1 and the same for Experiment 2. All observers had normal color vision verified with the Farnsworth-Munsell 100-Hue test; Ishihara pseudoisochromatic plates (24-plates edition, 1964); Rayleigh and Moreland anomaloscopy; and luminance matching (Interzeag Color Vision meter 712, Schlieren, Switzerland). All had normal or corrected-to-normal visual acuity. All observers were unaware of the purpose of the experiment.

Results

The frequency of “illuminant-change” responses in each condition was calculated as a function of the chromaticity of the local illuminant in the CIE 1976 (u' , v') chromaticity diagram. The frequency plots were smoothed by a two-

dimensional locally weighted quadratic regression (“loess” [11]), and the (u' , v') coordinates of the maximum of each distribution obtained numerically from the loess analysis. Deviations from the ideal setting, which corresponds to the second global illuminant and represents perfect color constancy, are listed in Tables 1 and 2.

Table 1. Errors in CIE 1976 (u' , v') coordinates of observers' most frequent “illuminant-change” responses with the urban scene of Experiment 1 under two different illuminants. Entries show means (and SEs) over the five initial test-surface colors.

	sky color shifts			
	neutral	blue	red	more red
u'	0.0041 (0.0009)	0.0030 (0.0003)	0.0031 (0.0010)	0.0046 (0.0008)
v'	0.0064 (0.0018)	0.0058 (0.0017)	0.0059 (0.0010)	0.0048 (0.0006)

Table 2. As for Table 1, but with the rural scene of Experiment 2.

	sphere-illuminant color shifts			
	neutral	red	more red	control
u'	0.0043 (0.0013)	0.0022 (0.0019)	0.0027 (0.0029)	0.0023 (0.0026)
v'	0.0074 (0.0046)	0.0062 (0.0023)	0.0066 (0.0051)	0.0069 (0.0052)

The results of the MANOVA applied to the (u' , v') coordinates of the most frequent “illuminant-change” responses were as follows. For Experiment 1 with the urban scene, there was no significant effect of visible sky color ($F(6, 30) = 0.73, p > 0.5$), although the effect of initial test color approached significance ($F(8, 40) = 2.1, p = 0.06$). For Experiment 2 with the rural scene, there was no significant effect of sphere-illuminant ($F(6, 30) = 0.37, p > 0.5$) nor initial test color ($F(8, 40) = 0.89, p > 0.5$).

In terms of a standard color-constancy index [12], for Experiment 1 with the urban scene, the mean index was 0.71, and for Experiment 2 with the rural scene, it was 0.70. Details are given elsewhere [7].

Discussion

Shifting the color of a directly visible sky along the daylight locus had no reliable effect on observers' surface-color matches; nor did shifting the color of the light from the sun and sky illuminating a visible sphere in the scene; nor did removing the sky from the field of view. Some biasing effect of the color of the sky or sphere on matches might have been expected; for, in addition to each being an explicit cue to the illuminant, each was also the brightest region of the scene, and, as such, its color should have influenced observers' general illuminant estimates [13]. It might be argued that other surfaces in the scenes were used for this purpose, but in Fig. 1 (c) and (d) the only surface with a diffuse highlight was in fact the sphere [14]. It might also be argued that observers assumed the illuminant was non-uniform, so that changes of its color in one location were unrelated to those in another, but again in Fig. 1 (c) and (d), both the sphere and the test surface were seen by the observer to be directly illuminated by the same source, namely the sun.

The implication of this analysis—that surface-color judgments do not depend on explicit knowledge of illuminant color—is consistent with a different experiment with non-natural scenes [5]. In that experiment, observers made asymmetric color matches between simultaneously presented Mondrian-like patterns of colored papers under different daylights. The patterns had either 49 surfaces or a minimal two surfaces, too few for an accurate estimate of the illuminant to be formed. Yet the degree of color constancy was almost exactly the same with the different numbers of surfaces.

A possible explanation of the insensitivity of surface-color judgments to information about the illuminant on a scene is that observers use “relational color constancy” [15]. This property refers to the constancy of perceived color relations between surfaces under different illuminants, as distinct from color constancy, which refers to the constancy of perceived colors of surfaces. Thus, in discriminating between illuminant and material changes in scenes, observers simply compare how the color of the test surface relates to the color of one or more other surfaces in the scene, or, indeed, to the scene as a whole, first under the one illuminant and then under the other. Knowledge of the illuminant itself is unnecessary, and manipulating an explicit cue to the color of the illuminant should leave performance unaffected, precisely as was found here.

Acknowledgements

Supported by the EPSRC (grant nos. GR/R39412/01 and EP/B000257/1).

References

- [1] M. D'Zmura and G. Iverson, Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces, *J. Opt. Soc. Am. A - Opt. Image Sci. Vis.*, 10, 2148 (1993).
- [2] M. D'Zmura and G. Iverson, Color constancy. II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces, *J. Opt. Soc. Am. A - Opt. Image Sci. Vis.*, 10, 2166 (1993).
- [3] M. D'Zmura and G. Iverson, Color constancy. III. General linear recovery of spectral descriptions for lights and surfaces, *J. Opt. Soc. Am. A - Opt. Image Sci. Vis.*, 11, 2389 (1994).
- [4] L. T. Maloney, Physics-based approaches to modeling surface color perception in Color Vision: From Genes to Perception. Cambridge University Press, Cambridge 1999, pp. 387.
- [5] K. Amano, D. H. Foster, and S. M. Nascimento, Minimalist surface-colour matching, *Perception*, 34, 1009 (2005).
- [6] D. H. Foster, S. M. C. Nascimento, and K. Amano, Information limits on neural identification of colored surfaces in natural scenes, *Visual Neurosci.*, 21, 331 (2004).
- [7] K. Amano, D. H. Foster, and S. M. Nascimento, Color constancy in natural scenes with and without an explicit illuminant cue, *Visual Neurosci.*, in press (2006).
- [8] D. H. Foster, K. Amano, and S. M. C. Nascimento, Colour constancy from temporal cues: better matches with less variability under fast illuminant changes, *Vision Res.*, 41, 285 (2001).
- [9] D. B. Judd, D. L. MacAdam, and G. Wyszecki, Spectral distribution of typical daylight as a function of correlated color temperature, *J. Opt. Soc. Am.*, 54, 1031 (1964).
- [10] D. H. Foster, K. Amano, and S. M. Nascimento, Color constancy in natural scenes explained by global image statistics., *Visual Neurosci.*, in press (2006).
- [11] W. S. Cleveland and S. J. Devlin, Locally weighted regression: an approach to regression analysis by local fitting, *J. Am. Statist. Assoc.*, 83, 596 (1988).

- [12] L. E. Arend, Jr, A. Reeves, J. Schirillo, and R. Goldstein, Simultaneous color constancy: papers with diverse Munsell values, *J. Opt. Soc. Am. A - Opt. Image Sci. Vis.*, 8, 661 (1991).
- [13] K. J. Linnell and D. H. Foster, Scene articulation: dependence of illuminant estimates on number of surfaces, *Perception*, 31, 151 (2002).
- [14] J. N. Yang and L. T. Maloney, Illuminant cues in surface color perception: tests of three candidate cues, *Vision Res.*, 41, 2581 (2001).
- [15] D. H. Foster and S. M. C. Nascimento, Relational colour constancy from invariant cone-excitation ratios, *Proc. R. Soc. Lond. Ser. B - Biol. Sci.*, 257, 115 (1994).

Author Biography

Kinjiro Amano received his Ph.D in 1998 from Tokyo Institute of Technology, Japan. He has worked at Aston University, UMIST, and is currently a research associate at the University of Manchester. His research work has concentrated on using visual psychophysics to investigate human color perception, in particular color constancy and color memory.