Estimating the information available from colored surfaces in natural scenes

Iván Marín-Franch and David H. Foster; Sensing, Imaging, and Signal Processing Group, School of Electrical and Electronic Engineering, University of Manchester; Manchester, UK

Abstract

The information preserved in identifying surfaces solely by their color can be quantified by measures defined by Shannon, including capacity and mutual information. The aim of this study was (1) to determine whether the capacity of an additive Gaussian channel provides an asymptotic upper bound to mutual information estimated for natural scenes under different illuminants, and (2) to explore the effect of different color representations on mutual-information estimates.

Introduction

Color is an imperfect code for representing surfaces in a scene. This is because the number of degrees of freedom in a sensor system, three for the cones in the normal human eye or a typical camera, is smaller than the number of degrees of freedom needed to specify different spectra [1]–[4]. Surfaces that match under one light therefore need not match under another: the phenomenon of metamerism. As a consequence, identifying surfaces on the basis of their color will lead to errors when the illuminant on scene changes. Such errors represent a loss of information. How much information, then, is preserved when surfaces are coded solely by their color?

One way to address this question is to estimate a quantity such as Shannon's mutual information [5, 6]. But making a direct estimate based on the probabilities involved leads to difficulties when those probabilities are small [7]. Instead, an asymptotic upper bound to the mutual information may be obtained from an analysis of the capacity of an additive Gaussian channel [8, 9]. The objective of the present work was to test whether capacity is an asymptotic upper bound to mutual information when estimated with samples of increasing size from natural scenes, and to explore how different color spaces [10, 11], color-difference formulae [12], and spectral sharpening [13] affect mutual-information estimates.

Methods

Hyperspectral Images

Scene reflectances were drawn from a set of eight hyperspectral images [8] (three of which are shown in Fig. 1). These images were from rural and urban areas in the Minho region of Portugal. The size of the images was approximately 1344×1024 pixels, and spectra at each pixel were defined at 10-nm wavelength intervals over 400–720 nm. Further technical details are available elsewhere [14, 8, 9].

Representation of Scenes

In computational simulations, scenes were illuminated successively by daylights of correlated color temperatures of 4000 K, 6500 K and 25000 K. For each of the three illuminants, the spectrum of the reflected light at each pixel was converted to tristimulus values and then to corresponding values (X_c, Y_c, Z_c), calculated with the CMCCAT2000 chromatic-adaptation transform [10] under the assumption of full adaptation and with D65 as reference. A version of CMCCAT2000 with sharp chromatic-adaptation transform [13, 15] was also used. The tristimulus values (X_c, Y_c, Z_c) were transformed to CIELAB values (L^*, a^*, b^*) with D65 as reference. Since CIELAB space is well known to be perceptually non-uniform [11], the color-difference formulae CIEDE2000 and CMCDE [12, 17] were each used to evaluate the differences in color-code values, in addition to the Euclidean distance.

Information-theoretic Measures and Estimates

For a particular scene, index the pixels in image 1 of the scene under illuminant e_1 by variable X and in image 2 of the scene under illuminant e_2 by variable Y. Suppose that in some task the probability of a particular pixel x in image 1 being chosen is $p(x) = P\{X = x\}$ and of a particular pixel y in image 2 being chosen is $p(y) = P\{Y = y\}$, where there are N pixels in each image. The degree of uncertainty associated with each image can be quantified by the entropy [5, 6] defined as follows:

$$H(X) = -\sum_{x=1}^{N} p(x) \log p(x), \tag{1}$$

with a similar expression for $H(Y)$. If the conditional probability $p(x|y) = P\{X = x|Y = y\}$ is known, then the conditional entropy is given by the expression:

$$H(X|Y) = -\sum_{y=1}^{N} p(y) \sum_{x=1}^{N} p(x|y) \log p(x|y). \tag{2}$$

This quantity represents the uncertainty about image 1 given image 2. The mutual information can then be expressed as the difference of the two entropies:

$$I(X;Y) = H(X) - H(X|Y). \tag{3}$$

Mutual information represents the reduction in uncertainty about image 1 given image 2. If the basis of the logarithm is 2, then the mutual information is expressed in bits (the convention adopted here).

As noted earlier, estimating $I(X;Y)$ directly from the probabilities leads to difficulties [7]. In principle, each pixel y under illuminant e_2 may be coded with value (L_2^*, a_2^*, b_2^*), say, and compared with each pixel x under illuminant e_1 coded with value (L_1^*, a_1^*, b_1^*) and then matched according the closest code value. Other color codes based on tristimulus values (X_c, Y_c, Z_c) or transformed CIELAB values (L', C', H') [12], with either CIEDE2000 or CMCDE color-difference formulae, can be used.

An alternative approach is to consider the distributions of the color-code values (L^*, a^*, b^*) under the two illuminants. An asymptotic upper bound C on the mutual information $I(X;Y)$, for large enough N, can be estimated from an analysis of the capacity of an additive Gaussian channel [6] by considering the differences in code values under the two illuminants as noise. The

Figure 1. Examples of pictures obtained with a hyperspectral imaging system [8].

capacity of this Gaussian channel has a formulation in terms of the covariance matrix of the code values of the scene under e_1 and covariance matrix of the noise [9]. Thus, let Σ_1 be the covariance matrix of the code values (L_1^*, a_1^*, b_1^*), and Σ_Δ the covariance matrix of the code value differences $(\Delta L^*, \Delta a^*, \Delta b^*)$. Then the capacity of the channel is given by the following:

$$C = \frac{1}{2} \log \left(\frac{|\Sigma_1 + \Sigma_\Delta|}{|\Sigma_\Delta|} \right). \tag{4}$$

Neither code values (L^*, a^*, b^*) nor the differences $(\Delta L^*, \Delta a^*, \Delta b^*)$ are exactly normally distributed. Even so, it may still be shown that (4) cannot be exceeded with normally distributed code values and nearest-neighbor identification [16].

Another approach to the estimation of $I(X;Y)$ is to calculate the number of pixels likely to be in error in each match [9]. The idea is to take a random sample, say n, from image 2 and count for each pixel x in that sample how many pixels in image 1 are within the tolerance defined by the error in matching, i.e. the number of pixels incorrectly matched, or number of errors, for x. For two different pixels, the error in matching is, in general, different, and, even if they are equal, the number of pixels incorrectly matched might differ. Since the pixels in image 2 are chosen randomly, the errors are also random, and the probability of having a particular number of errors can then be estimated. Once the probabilities of error are estimated, the conditional entropy (2) follows. If the sample is random and uniform, the entropy (1) is equal to $\log n$, and the mutual-information estimate follows from (3).

To take into account the variance-covariance structure of the differences in color-code values due to illuminant changes, the Mahalanobis (statistical) distance was used to quantify the goodness of the match. For CIEDE2000 and CMCDE color-difference formulae, in which the distance is already determined [12, 17], mutual information was not estimated.

Results and Comment

Capacity and mutual information were estimated using both CIE 1931 2° and CIE 1964 10° standard observers. Despite variations in the actual values, the differences between capacity and mutual-information estimates were about the same. The results shown here correspond to the CIE 1964 10° standard observer, chosen for its compatibility with the CIEDE2000 color-difference formula [17]

Channel Capacity Estimation

Tables 1–3 show the mean channel capacity C (and standard deviation) estimated from (4) for the eight scenes, with three different color codes: normalized and unnormalized (X_c, Y_c, Z_c) (Table 1); normalized and unnormalized CIELAB (L^*, a^*, b^*) (Table 2); and transformed CIELAB (L', C', H') [12] (Table 3), this last code using CIEDE2000 and CMCDE color-difference formulae. The illuminant changes were 4000 K to 6500 K, 6500 K to 25000 K, and 4000 K to 25000 K. The tables are divided into upper and lower sections: the upper for complete adaptation under CMCCAT2000, and the lower the same but with a sharp chromatic-adaptation transform. The normalized data in the last two columns of Tables 1 and 2 refer to a repetition of (X_c, Y_c, Z_c) and (L^*, a^*, b^*) codes but after applying an empirical transformation that forced the distributions to be Gaus-

Table 1. Mean capacity C (and SD) in bits for eight scenes under three different illuminant changes and (X_c, Y_c, Z_c) color code. Upper and lower sections for unsharp and sharp codes.

		Unnormed	Normed
Unsharp	4000 K – 6500 K	12.42 (0.53)	11.42 (0.79)
	6500 K – 25000 K	12.33 (0.70)	11.38 (0.88)
	4000 K – 25000 K	09.57 (0.55)	08.64 (0.82)
Sharp	4000 K – 6500 K	15.86 (0.72)	14.28 (1.11)
	6500 K – 25000 K	15.03 (0.70)	13.62 (1.00)
	4000 K – 25000 K	12.51 (0.64)	11.08 (1.00)

Table 2. Details as for Table 1 but for (L^*, a^*, b^*) color codes.

		Unnormed	Normed[a]
Unsharp	4000 K – 6500 K	12.32 (0.55)	11.90 (0.45)
	6500 K – 25000 K	12.27 (0.57)	11.57 (0.42)
	4000 K – 25000 K	09.47 (0.52)	09.00 (0.39)
Sharp	4000 K – 6500 K	15.83 (0.76)	14.79 (0.89)
	6500 K – 25000 K	14.98 (0.78)	13.90 (0.63)
	4000 K – 25000 K	12.46 (0.65)	11.54 (0.65)

[a] Scene (c) of Fig. 1 cropped to remove dark areas

Table 3. Details as for table 1 but for $(L'C'H')$ color codes and both CIEDE2000 and CMCDE color-difference formulae.

		CIEDE2000	CMCDE
Unsharp	4000 K – 6500 K	11.24 (1.31)	11.53 (1.33)
	6500 K – 25000 K	10.66 (1.34)	11.09 (1.36)
	4000 K – 25000 K	08.60 (0.88)	08.93 (0.90)
Sharp	4000 K – 6500 K	13.63 (1.64)	13.94 (1.67)
	6500 K – 25000 K	12.53 (1.58)	13.04 (1.65)
	4000 K – 25000 K	10.58 (1.25)	10.96 (1.31)

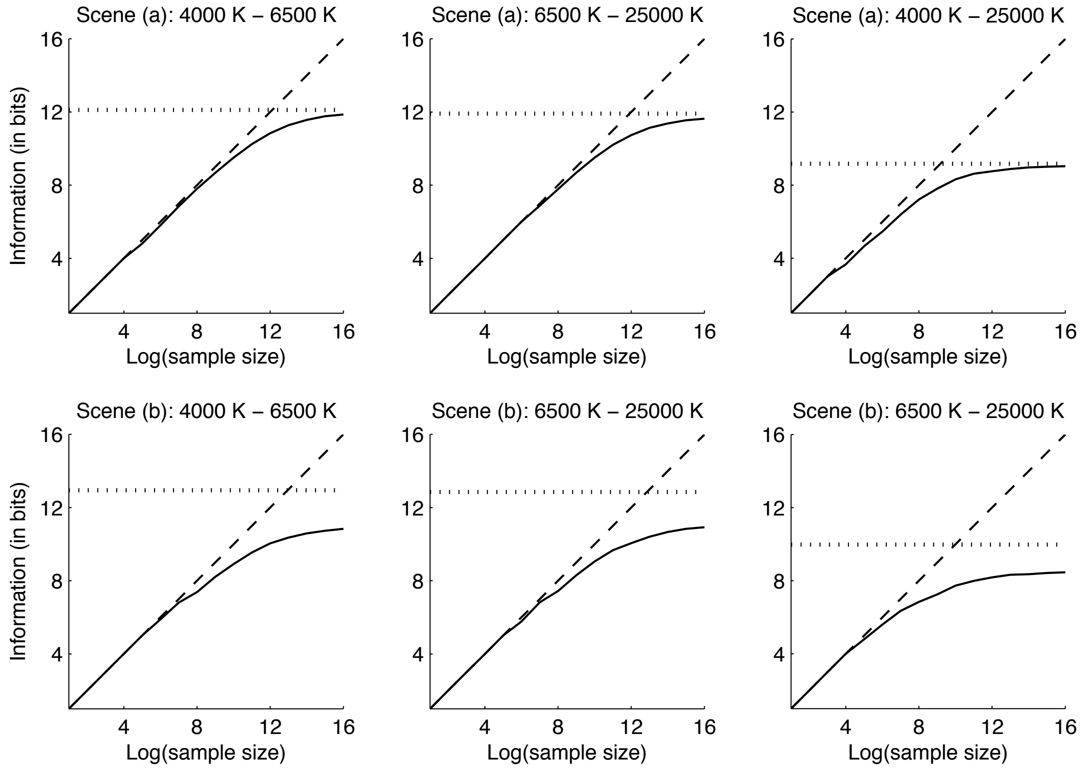

Figure 2. *Mutual information as a function of sample size (solid curve) and channel capacity (dotted line) for scenes (a) and (b) of Fig. 1 for* $(L^*,a^*,b^*)'$ *color code. The diagonal dashed line show the maximum information as function of sample size.*

sian with the same mean and standard deviation as the original. This transformation was introduced because the distributions of (X_c,Y_c,Z_c), (L^*,a^*,b^*) and (L',C',H') code values are not normal and, as noted earlier, the calculation of capacity depends on the channel being Gaussian. Scene (c) of Fig. 1 yielded an abnormally low capacity estimate for the (L^*,a^*,b^*) color code after normalization owing to the presence of dark areas, which produced lightness differences ΔL^* that were far from normal. This scene was cropped to remove these areas.

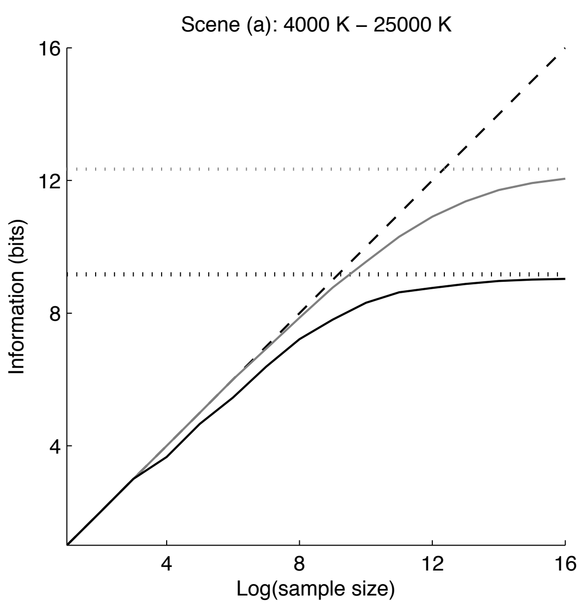

Figure 3. *Effect of a sharp chromatic-adaptation transform on capacity and mutual-information estimates. Details as for Fig. 2. Solid curve and line are for unsharp data and gray curve and line for sharp data.*

Capacity estimates for (X_c,Y_c,Z_c) and (L^*,a^*,b^*) color codes were about 1 bit larger than for (L',C',H') codes. For sharp chromatic-adaptation transforms, the differences were even larger (about 2 bits). In contrast, estimates for the normalized (X_c,Y_c,Z_c) and (L^*,a^*,b^*) color codes were closer (lower than 0.4 and 0.8 bits on average for unsharp and sharp chromatic-adaptation transforms, respectively). On average, spectral sharpening increased capacity estimates by 3.0 and 3.1 bits for the unnormalized (X_c,Y_c,Z_c) and (L^*,a^*,b^*) color codes, respectively; by 2.5 and 2.6 bits for the normalized (X_c,Y_c,Z_c) and (L^*,a^*,b^*), respectively; and by 2.1 bits for the (L',C',H') codes with both color-difference formulae. For any particular illuminant change, capacity estimates were reasonably stable over color codes: the standard deviation varied from 0.4 to 1.4 bits with CMCCAT2000 and from 0.6 to 1.7 bits with a sharp chromatic-adaptation transform. Estimates for (L',C',H') with both color-difference formulae were very similar across illuminant changes, and closer to the normalized (X_c,Y_c,Z_c) and (L^*,a^*,b^*) values than to the unnormalized values.

Mutual Information Estimation

As with capacity estimates, for any particular illuminant change, mutual-information estimates remained reasonably stable over color codes (for a sample size of 2^{14}). Mutual-information estimates also increased with sharp chromatic-adaptation transforms. Normalization had a small effect on mutual-information estimates for (X_c,Y_c,Z_c) and (L^*,a^*,b^*)

codes, but in contrast to capacity estimates, it led to an increase rather than a decrease in the estimate.

Fig. 2 shows mutual-information estimates (solid curve) for (L^{*}, a^{*}, b^{*}) color codes as a function of sample size (here extended to 2^{16}), and the corresponding capacity C (horizontal dotted line). Scene labels refer to Fig. 1. For image (b) mutual-information estimates were lower than capacity estimates, but this difference was about the same for all illumination changes. Similar results were obtained for the other scenes with $(X_{\mathrm{c}}, Y_{\mathrm{c}}, Z_{\mathrm{c}})$. Fig. 3 shows mutual-information estimates for sharp and unsharp chromatic-adaptation transforms. The dependence on sample size in the two conditions is similar, but with the sharp transform the gradient of the information function decreases more slowly with sample size than with the unsharp transform, consistent with their different asymptotic limits. Similar plots were obtained for other scenes and illumination changes. Spectral sharpening seems to have the same effect on capacity as on mutual-information estimates.

Conclusions

The existence of metamerism implies that information is lost when identifying surfaces in a scene solely on the basis of their color. As shown here, the capacity of an additive Gaussian-channel model provides an approximate asymptotic upper bound on the estimate of the information preserved, the value depending on the particular scene, illuminant change, and color representation. A potential application of this approach is in providing an objective measure of the efficiency of different coding strategies in the presence of unreliable image data.

Acknowledgments

We thank K. Żychaluk and S. Panzeri for useful discussions, and S. M. C. Nascimento and K. Amano for assistance in providing the hyperspectral images. This work was supported by the Engineering and Physical Sciences Research Council (grant no. EP/B000257/1).

References

[1] L. T. Maloney, Evaluation of linear models of surface spectral reflectance with small numbers of parameters, J. Opt. Soc. Am. A 3, 1673 (1986).

[2] J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, Characteristic spectra of munsell colors, J. Opt. Soc. Am. A 6, 318 (1989).

[3] S. M. C. Nascimento, D. H. Foster, and K. Amano, Psychophysical estimates of the number of spectral-reflectance basis functions needed to reproduce natural scenes, J. Opt. Soc. Am. A 22, 1017 (2005).

[4] E. K. Oxtoby and D. H. Foster, Perceptual limits on low-dimensional models of munsell reflectance spectra. Perception 34, 961 (2005).

[5] C. E. Shannon, A mathematical theory of communication., Bell Syst. Tech. J. 27, 379 (1948).

[6] T. M. Cover and J. A. Thomas, Elements of Information Theory, John Wiley & Sons, Inc. New York (1991).

[7] D. R. Brillinger, Second-order moments and mutual information in the analysis of time series, in Recent Advances in Statistical Methods, pg. 64, Imperial College Press (2001).

[8] 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).

[9] D. H. Foster, S. M. C. Nascimento, and K. Amano, Information limits on identification of natural surfaces by apparent colour, Perception 34, 1001 (2005).

[10] C. J. Li, M. R. Luo, B. Rigg, and R. W. G. Hunt, CMC 2000 Chromatic Adaptation Transform: CMCCAT2000, Color Res. Appl. 27, 49 (2002).

[11] M. Fairchild, Color appearance models, Wiley-IS&T, Chichester, UK (2005).

[12] M. R. Luo, G. Cui, and B. Rigg, The development of the CIE 2000 colour-difference formula: CIEDE2000, Color Res. Appl. 26, 340 (2001).

[13] G. D. Finlayson, M. S. Drew, and B. V. Funt, Spectral sharpening: sensor transformations for improved color constancy, J. Opt. Soc. Am. A 11, 1553 (1994).

[14] S. M. C. Nascimento, F. P. Ferreira, and D. H. Foster, Statistics of spatial cone-excitation ratios in natural scenes, J. Opt. Soc. Am. A 19, 1484 (2002).

[15] S. Susstrunk, J. Holm, and G. D. Finlayson, Chromatic adaptation performance of different RGB sensors, Proc. SPIE, pg. 172 (2001).

[16] A. Lapidoth, Nearest neighbor decoding for additive non-gaussian noise channels, IEEE Trans. Inf. Theory 42, 1520 (1996).

[17] CIE, Colorimetry, 3rd Edition, CIE Publication 15:2004. CIE Central Bureau, Vienna (2004).

Author Biography

Iván Marín-Franch received his B.Sc. in Statistics in 1998 and M.Sc. in Statistical Science and Techniques in 2000 from the University of Granada. In 2005, he began his research for a Ph.D. in Sensing, Imaging, and Signal Processing in the Faculty of Engineering and Physical Sciences of the University of Manchester.