ELSEVIER

VISION
RESEARCH

CrossMark

Time-lapse ratios of cone excitations in natural scenes

David H. Foster [a,*], Kinjiro Amano [a], Sérgio M.C. Nascimento [b]

[a] School of Electrical and Electronic Engineering, University of Manchester, Manchester M13 9PL, UK
[b] Centre of Physics, University of Minho, Campus de Gualtar, 4710-057 Braga, Portugal

ABSTRACT

The illumination in natural environments varies through the day. Stable inferences about surface color might be supported by spatial ratios of cone excitations from the reflected light, but their invariance has been quantified only for global changes in illuminant spectrum. The aim here was to test their invariance under natural changes in both illumination spectrum and geometry, especially in the distribution of shadows. Time-lapse hyperspectral radiance images were acquired from five outdoor vegetated and non-vegetated scenes. From each scene, 10,000 pairs of points were sampled randomly and ratios measured across time. Mean relative deviations in ratios were generally large, but when sampling was limited to short distances or moderate time intervals, they fell below the level for detecting violations in ratio invariance. When illumination changes with uneven geometry were excluded, they fell further, to levels obtained with global changes in illuminant spectrum alone. Within sampling constraints, ratios of cone excitations, and also of opponent-color combinations, provide an approximately invariant signal for stable surface-color inferences, despite spectral and geometric variations in scene illumination.

1. Introduction

In natural environments, the light from the sun and sky varies continuously over the course of the day. As a result, the reflected light provides a constantly changing signal to the eye. As Larry Arend remarked,

> The retinal signal is jointly determined by some combination of distal optical variables, some of which (the ones we wish to recover) are intrinsic color properties of the surface and some of which are contingent, i.e., accidents of the transient optical environment in which the surface is being viewed.
>
> [Arend (2001), p. 392]

Yet despite this changing signal, our perception of natural scenes is stable. What, then, underpins this stability, in particular, the stability of perceived surface colors under different illuminations? As with geometrical shape perception under changes in object pose, taking ratios of signals can sometimes deliver an invariant property (e.g. Kent & Mardia, 2012; Maybank, 1995; Moons, Pauwels, Van Gool, & Oosterlinck, 1995). A retinally derived signal that offers an invariance to changes in illumination is the spatial ratio of cone excitations generated in response to

light reflected from different surfaces (Foster & Nascimento, 1994). From computational simulations with natural scenes (Nascimento, Ferreira, & Foster, 2002), it is known that spatial ratios remain approximately invariant under changes in the illuminant.

This invariance property has long been assumed, explicitly or implicitly, in chromatic-adaptation models of the von Kries type (Brill, 2008; van Trigt, 2007; von Kries, 1902, 1905) and in the Retinex models due to Land (Funt, Ciurea, & McCann, 2004; Land, 1983; Land & McCann, 1971). Spatial ratios have also been used to explain judgments about surface color in displays of Munsell colored papers under different illuminants (Amano, Foster, & Nascimento, 2005), even when chromatic adaptation does not eliminate differences in color appearance (Reeves, Amano, & Foster, 2008). Moreover, in operationally oriented tasks, spatial ratios seem to be the kinds of signals preferred by observers for making discriminations between illuminant changes and simulated changes in surface reflectances (Nascimento & Foster, 1997).

Nonetheless the invariance of spatial ratios in both natural scenes and Munsell papers has been quantified only with simulated global changes in illumination spectrum (Nascimento, Ferreira, & Foster, 2002). This limitation is immaterial providing that interest is in spatially uniform illuminants (e.g. Foster, Amano, Nascimento, & Foster, 2006). But it does become relevant in the transient optical environment of the natural world where changes in the spectrum

of the illumination are accompanied by changes in its geometry. Those changes may be global, due to movement of cloud and variations in atmospheric scatter, or local, due to changes in mutual illumination (Bloj, Kersten, & Hurlbert, 1999; Funt & Drew, 1993) and in shadows (Giles, 2001; Zhou, Huang, Troy, & Cadenasso, 2009), which may be attached, depending on the angle of incidence of the illumination and viewing direction, or unattached or cast, depending on other objects in the scene. Because of the general complexity of natural scenes (Arend, 2001; Endler, 1993), changes in mutual illumination and in the distribution of shadows probably represent the greatest challenge to the automatic extraction of surface-color attributes.

Most studies of the role of illumination geometry in perception have been concerned with object shape or surface perception (e.g. Bloj, Kersten, & Hurlbert, 1999; Castiello, 2001; Cavanagh & Leclerc, 1989; Elder, Trithart, Pintilie, & MacLean, 2004; Lee & Brainard, 2014; Leek, Davitt, & Cristino, 2015; Mamassian, Knill, & Kersten, 1998; Ripamonti et al., 2004; Schofield, Rock, & Georgeson, 2011; Tarr, Kersten, & Bülthoff, 1998; Wagemans, van Doorn, & Koenderink, 2010). But some studies have considered the distributional effects of shadows, for example, in matching and discrimination with regular geometrical arrays (e.g. Heckman, Muday, & Schirillo, 2005; Kingdom, Beauce, & Hunter, 2004); in detecting fruit in natural scenes (Lovell et al., 2005), and in relation to foraging behavior (Arnold & Chittka, 2012).

The effects of changing illumination geometry on reflected spectra can sometimes be unexpected. Fig. 1 shows sampled spectra and color images rendered from hyperspectral radiance images of a natural scene at different times of the day. The spectra at locations 1 and 2 in the foreground foliage are almost constant between 11:40 and 12:40 (graphs 1 and 2, left and center columns), yet appear to interchange as the direction of illumination modulates the detailed pattern of shadows between 12:40 and 16:37 (graphs 1 and 2, center and right columns). The spectrum at location 3 in the distant haze and woodland retains its profile but undergoes progressive compression as the amount of haze diminishes (graph 3, left, center, and right columns). By contrast, over the same period, the profile of the spectrum at location 4 from the midfield red roof remains almost constant (graph 4, left, center, and right columns).

The aim of this study was to test how well ratios of cone excitations sampled from natural scenes are preserved under natural changes in illumination, especially in the distribution of shadows, through the course of the day. Time-lapse sequences of hyperspectral radiance images were acquired from five outdoor vegetated and nonvegetated scenes. From each scene, 10,000 pairs of points were sampled randomly and ratios of cone excitations and of other postreceptoral combinations of cone excitations were measured across time. Estimates of the mean relative deviations in ratios were found to be generally large. But when sampling was limited to short distances or to moderate time intervals, they fell below the criterion level for detecting violations of ratio invariance. Additionally, when uneven changes in illumination geometry were excluded, for example, when sample points were first in direct sunlight and then partially in shade, they fell further, to levels obtained with simulated global changes in illumination spectrum. Within sampling constraints, ratios provide an approximately invariant signal for stable surface-color inferences.

Partial reports of these findings were presented at the Vision Sciences Society 15th Annual Meeting (VSS 2015) and at MODVIS 2015, St. Pete Beach, FL, USA, 2015. Preliminary accounts of time-lapse changes in cone-excitation ratios were included in unpublished reports from two Master's projects (Dörrer & Newton, 2007; Raath & Woodward, 2008).
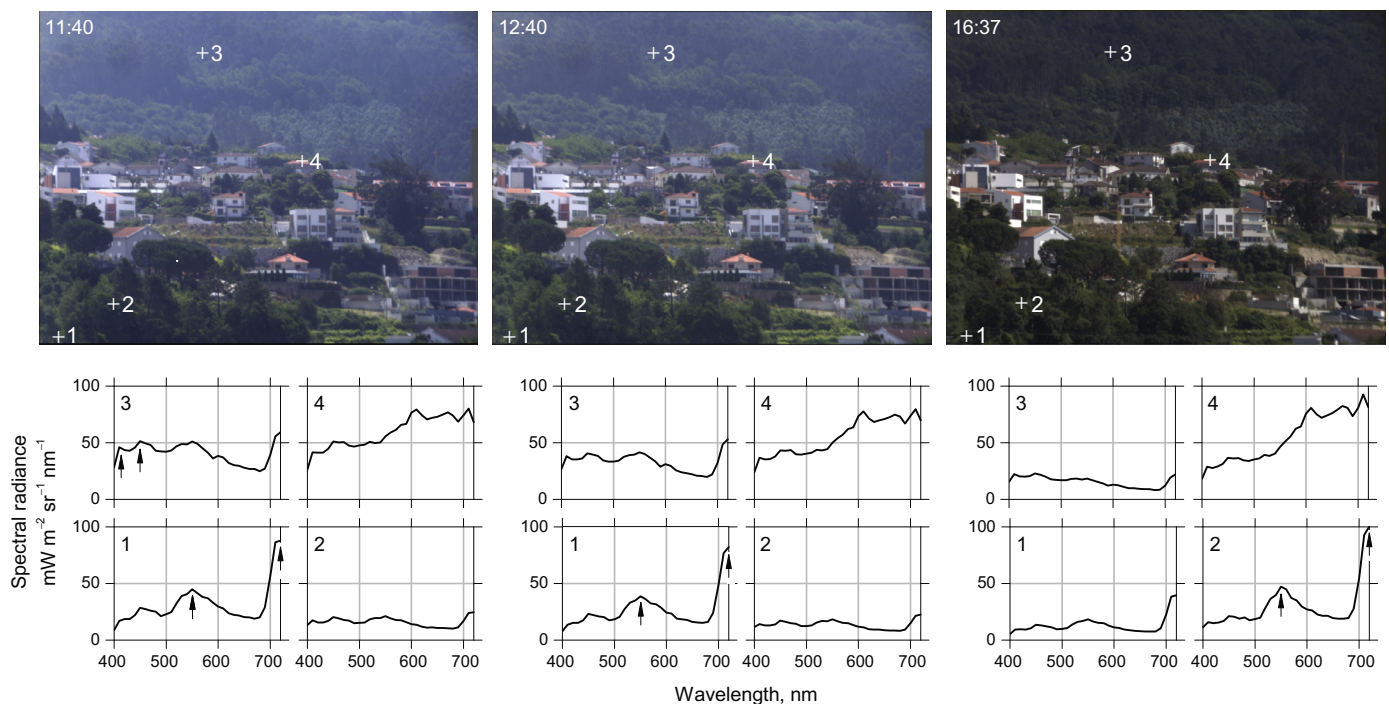
Fig. 1. Top row. Color images rendered from three hyperspectral radiance images of the parish of Nogueiró in the Minho region of Portugal. The images were acquired at 11:40, 12:40, and 16:37 on 9 June 2003. Bottom row. Plots of reflected radiance spectra at single-pixel locations marked by crosses in the images above. For the spectra at locations 1 and 2 in the foreground foliage, the maxima at 550 and 720 nm (arrowed, left, center, and right columns) are similar to those found in radiance estimates derived from laboratory spectral reflectance measurements (Carter & Knapp, 2001; Sims & Gamon, 2002). For the spectrum at location 3 in the distant woodland, the local maxima at 410 and 450 nm (arrowed, left column) were attributed to haze and are similar to those found in typical daylight spectra (Hernández-Andrés, Romero, & Nieves, 2001; Judd et al., 1964; see also Burton & Moorhead, 1987). Other aspects of the spectra are discussed in the text.

2. Methods

2.1. Image acquisition

The custom-built hyperspectral imaging system has been described elsewhere (Foster et al., 2006, Sect. 2). The essential details are reproduced here, with some additional calibration information. The system comprised a low-noise Peltier-cooled digital camera with CCD sensor (Hamamatsu, model C4742-95-12ER, Hamamatsu Photonics K. K., Japan) and a fast liquid–crystal tunable filter (VariSpec, model VS-VIS2-10-HC-35-SQ, Cambridge Research & Instrumentation, Inc., MA) mounted in front of the lens, together with an infrared blocking filter. The camera provided an effective image size of 1344 × 1024 pixels. The focal length was set typically to 75 mm and aperture to ƒ/16 or ƒ/22 to achieve a large depth of focus. The zoom was set to maximum. The system line-spread function was close to Gaussian with a standard deviation of approximately 1.3 pixels at 550 nm. The intensity response at each pixel, recorded with 12-bit precision, was linear over the entire dynamic range. The wavelength range of the tunable filter was 400–720 nm. Its bandwidth (FWHM) was 7 nm at 400 nm, 10 nm at 550 nm, and 16 nm at 720 nm. Its tuning accuracy was <1 nm over 400–660 nm and <2 nm over 670–720 nm, and it differed by <1 nm between the center and at the edges of the imaging field (see Ekpenyong (2013) and Pinto (2004) for further details). Root-mean-square errors in the recovery of radiance spectra were <2%, based on measurements of the recovery of spectral reflectance factors of colored samples (Foster et al., 2006; Nascimento, Ferreira, & Foster, 2002).

At each scene, the camera system was adjusted for position, viewing direction, focus, and aperture, all of which were then fixed. A small neutral (Munsell N7) flat surface was embedded in the scene to provide a spectral reference and its position and orientation adjusted in relation to the field of view. A small neutral (Munsell N7) glass or plastic sphere was also embedded in all but one scene to help quantify the effects of changing illumination geometry. Around the time of image acquisition, a horizontally oriented barium sulfate plug was introduced to enable estimates to be made of the illumination under direct sunlight and in the shade. The reflected spectrum was recorded with a telespectroradiometer (SpectraColorimeter, PR-650, Photo Research Inc., Chatsworth, CA), whose calibration was traceable to the National Physical Laboratory. The angular subtense of each scene at the camera was approximately 6.9° × 5.3°. Each pixel therefore represented the integrated optical image radiance over approximately 0.3 × 0.3 arcmin.

Each time-lapse sequence of hyperspectral images was obtained by making acquisitions at intervals of about one hour, sufficiently long to reveal changes in the distribution of illumination (see e.g. Fig. 1), although with some scenes acquisitions were made more frequently towards the end of the day. Each acquisition, which took a few minutes, comprised three steps. First, the camera exposure time required at each wavelength was determined by an automatic calibration routine that ensured the maximum sensor-element output was within 86–90% of its saturated value. Second, with exposure times defined, wavelength sequences of raw grayscale images were recorded over the range 400–720 nm at 10-nm intervals. The recording was repeated one or more times within each acquisition period to enable the signal-to-noise ratio (SNR) to be improved, as explained later. Third, the spectral radiance of the light reflected from the reference surface was recorded from the same viewing direction with the telespectroradiometer. Depending on the scene, acquisitions continued to be made over 9–10 h or until the level of natural light impeded recording.

2.2. Image processing

The raw grayscale images were first corrected by subtracting estimates of dark noise and stray light derived from earlier calibration measurements. Spatial non-uniformities (mainly off-axis vignetting) were corrected by dividing by an estimate of the flat-field image. Where recordings of wavelength sequences of images had been repeated (Section 2.1), the corrected images at each wavelength were registered over replicates by translation to compensate for any small differences in optical image position and then averaged. They were next registered over wavelength by uniform scaling and translation to compensate for any small differences in optical image size with wavelength (chromatic differences of magnification) and any small differences in optical image position. A few wavelength sequences of images were omitted because their SNRs at 400 nm made registration unreliable. The resulting hyperspectral image was then calibrated spectrally by applying a spatially uniform spectral correction factor such that the spectral radiance at the neutral reference surface in the image matched the spectral radiance recorded with the telespectroradiometer. Recall that tuning accuracy differed by <1 nm between the center and at the edges of the imaging field. The calibrated hyperspectral radiance images from each acquisition period were then assembled to make a time-lapse sequence. Finally, the sequence was registered over acquisition times by translation to compensate for any residual differences in optical image position. All image registrations were performed to subpixel accuracy with in-house software.

No assumptions were made about the spectral reflecting properties of individual surfaces within the scenes (apart from the reference surfaces) or the spectral properties of the differing illuminations on the scenes (cf. Foster et al., 2006; Nascimento, Ferreira, & Foster, 2002).

2.3. Scenes

Imaging took place in June and October of 2003 in the Minho region of Portugal. The five outdoor scenes chosen contained a mixture of woodland, herbaceous vegetation, barren land and rock, and rural and urban buildings. The scenes were the parish of Nogueiró, a building in Gualtar, a rock face in Sete Fontes, the village of Levada, and a terrace in Sameiro. Images were acquired under direct sunlight in a clear or almost-clear sky. If clouds were present, care was taken to avoid variations in illumination during the acquisition. Color images rendered from the hyperspectral images of each scene are shown in Fig. 2, together with the times and dates of acquisition. Over scenes, the correlated color temperature of the illumination from the sun and sky ranged from 4400 K to 8300 K in direct sunlight and from 6000 K to 22,600 K in the shade.

For the Nogueiró, Sete Fontes, and Levada scenes, three replicate images were available for averaging at each wavelength from each acquisition period (Sections 2.1 and 2.2). For the Gualtar scene, there were two. But for the Sameiro scene, imaging conditions proved more difficult than expected, with the result that no reliable replicates could be obtained. Although for the sake of completeness, data from unaveraged Sameiro hyperspectral images are included in an illustration of the effects of the spacing of sample points in scenes (Section 3.2), they were otherwise excluded from the subsequent analysis.

Data from these time-lapse sequences of hyperspectral radiance images have not been reported previously, except for one hyperspectral image of the 7–9 from each scene being used for spectral-reflectance estimates (e.g. Foster et al., 2006, Sect. 2).

Nogueiró

Gualtar

Sete Fontes

Levada

Sameiro
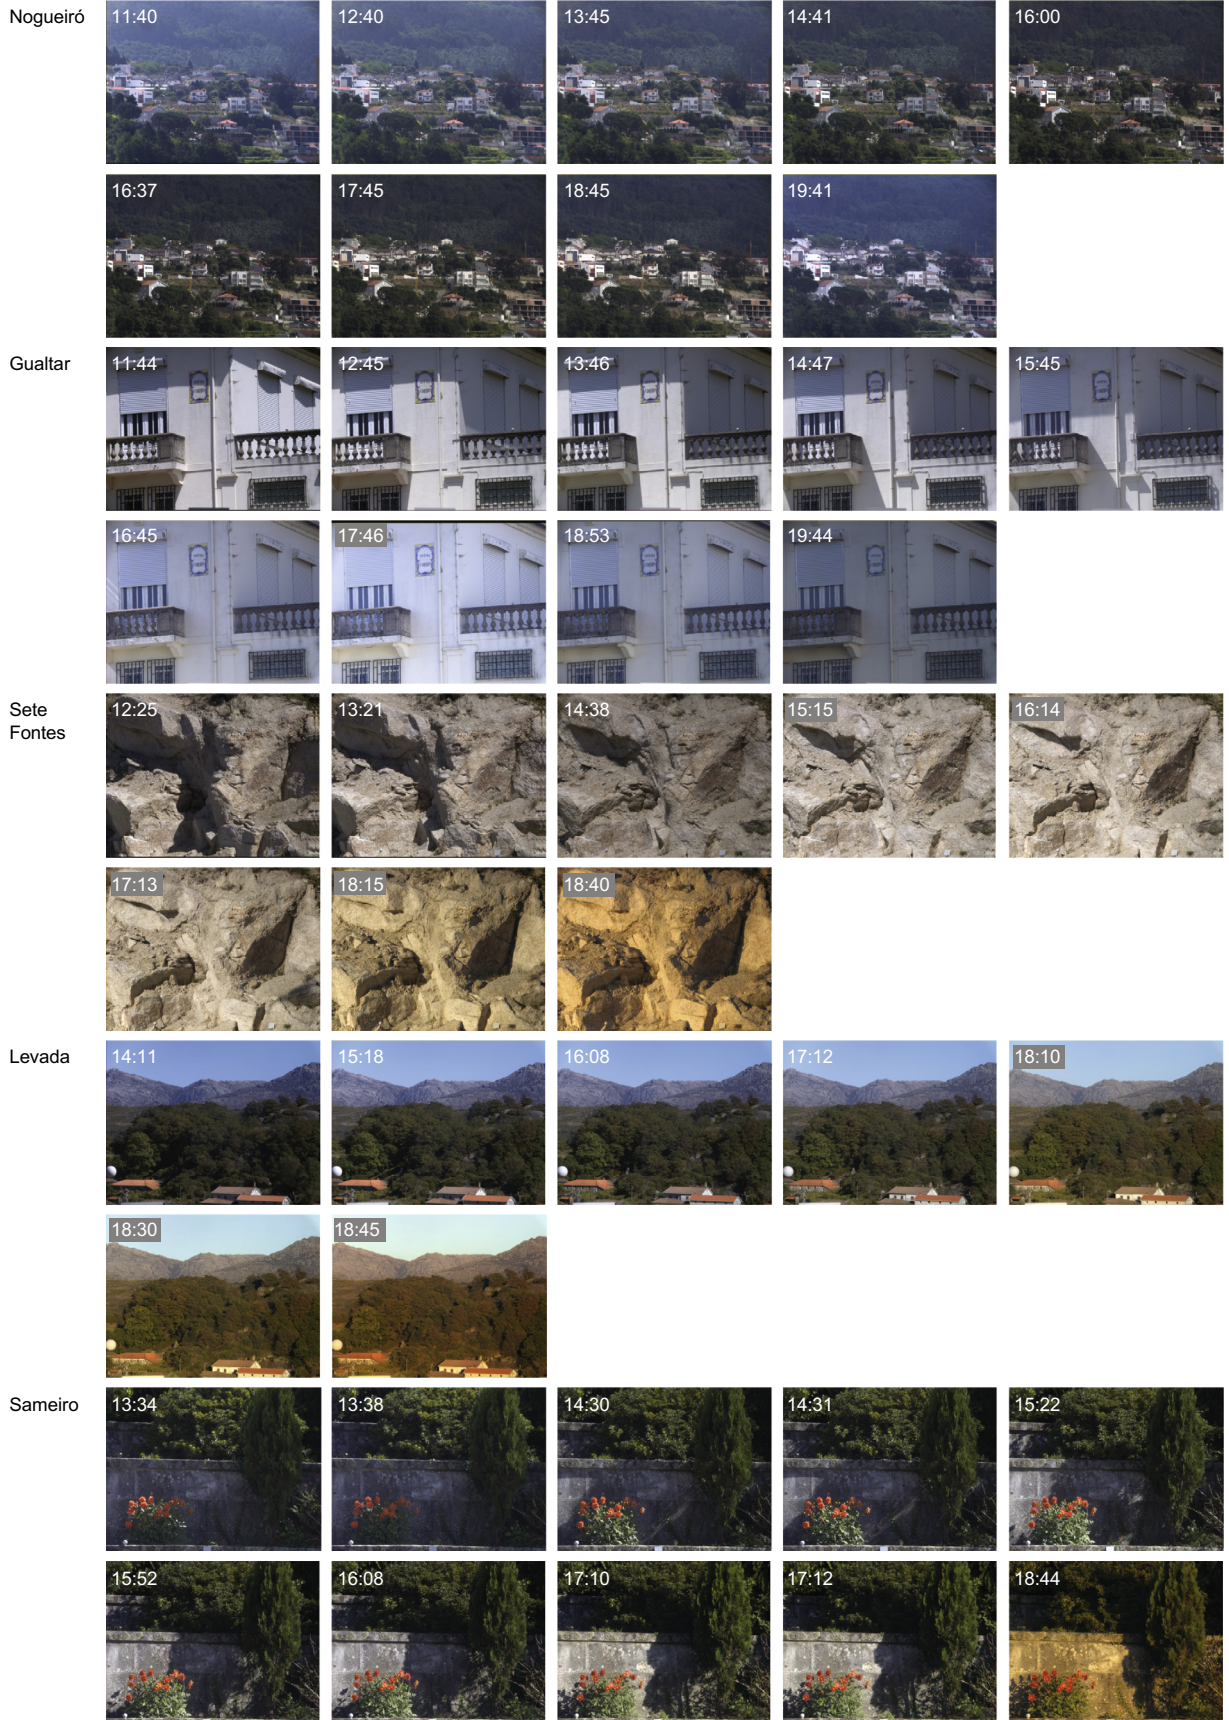
Fig. 2. Color images rendered from time-lapse sequences of hyperspectral radiance images. The images for all but the last scene were obtained by averaging multiple replicates from each acquisition period. The scenes were the parish of Nogueiró, a building in Gualtar, a rock face in Sete Fontes, the village of Levada, and a terrace in Sameiro, all in the Minho region of Portugal. The dates of image acquisition were 9 June 2003, 6 June 2003, 6 October 2003, 8 October 2003, and 7 October 2003, respectively. Acquisition times are shown in the top left of each color image. Image sizes were 1344 × 1024 pixels, and each scene subtended approximately 6.9° × 5.3° at the camera from a distance of approximately 1.3 km, 53 m, 29 m, 560 m, and 25 m, respectively. The Munsell reference surfaces can be seen as either narrow strips or square patches on the right-hand side of the Nogueiró scene and at the bottom of the Gualtar, Sete Fontes, Levada, and Sameiro scenes. Some of the color images have been brightened for the purposes of illustration. The thin black or colored edges present in some images are the result of multiple hyperspectral image registrations, mainly over time, and were excluded from the analysis.

2.4. Cone excitations and spatial ratios

Data values in each time-lapse sequence of hyperspectral images were indexed by acquisition time t, wavelength λ, and spatial coordinates (x, y). For each scene, a sample of $2N$ points (x_i, y_i) and (x_i', y_i'), with $i = 1, 2, \ldots, N$, was drawn randomly without replacement from the 1,376,256 (= 1344 × 1024) available.[1] Sampling was constrained so that the spacing of sample points, i.e. the distance between (x_i, y_i) and (x_i', y_i'), was drawn from the logarithmic scale 1, 2, 4, …, 256 pixels. As shown later, a logarithmic scale linearized deviations in ratios. Spacings greater than 256 pixels produced obvious sampling nonuniformities near the boundary of the 1344 × 1024 pixels array. The value of N was varied in pilot simulations and $N = 10,000$ was found to be sufficient for repeatable estimates of the mean (each estimated SE was of the order of 1% of the estimated mean).

Cone excitations were calculated in the usual way. Suppose that at each pair of points (x_i, y_i) and (x_i', y_i'), the spectral radiances at wavelength λ and time t were $E_t(\lambda; x_i, y_i)$ and $E_t(\lambda; x_i', y_i')$, respectively.[2] Let $L(\lambda)$, $M(\lambda)$, and $S(\lambda)$ be the long-, medium-, and short-wavelength-sensitive (L, M, and S) corneal cone spectral sensitivities, respectively. Then the corresponding cone excitations $l_i, m_i,$ and s_i at (x_i, y_i) are given by the integrals:

$$l_i = \int E_t(\lambda; x_i, y_i) L(\lambda) \, d\lambda,$$
$$m_i = \int E_t(\lambda; x_i, y_i) M(\lambda) \, d\lambda,$$
$$s_i = \int E_t(\lambda; x_i, y_i) S(\lambda) \, d\lambda,$$

and the corresponding cone excitations $l_i', m_i',$ and s_i' at (x_i', y_i') analogously.[3] Integrals were evaluated over the wavelength range 400–720 nm by summing at 10-nm intervals. The cone spectral sensitivities $L(\lambda)$, $M(\lambda)$, and $S(\lambda)$ were taken from the Stockman and Sharpe 2-degree cone fundamentals (Stockman & Sharpe, 2000; Stockman, Sharpe, & Fach, 1999).

For each cone class, let r_i be the ratio of excitations at (x_i, y_i) and (x_i', y_i'), e.g. for L cones, $r_i = l_i/l_i'$, with $i = 1, 2, \ldots, N$; and let $r_{i,1}$ and $r_{i,2}$ be the particular values of these ratios at times t_1 and t_2, respectively. It is helpful to normalize the differences $r_{i,1} - r_{i,2}$ by $r_{i,1}$ or $r_{i,2}$ or some combination of the two so as to stabilize the variance. The mean relative deviation MRD between ratios $r_{i,1}$ and $r_{i,2}$ over all sample pairs $i = 1, 2, \ldots, N$ was thus defined by:

$$\mathrm{MRD} = \frac{1}{N} \sum_{i=1}^{N} \frac{|r_{i,1} - r_{i,2}|}{(r_{i,1} + r_{i,2})/2}. \tag{1}$$

For comparison with an alternative definition of mean relative deviation, see Appendix A.

Changes of illumination geometry at (x_i, y_i) and (x_i', y_i') across times t_1 and t_2 were classified as even or uneven according to a test statistic derived from the reflected spectral radiances $E_1(\lambda; x_i, y_i)$ and $E_1(\lambda; x_i', y_i')$ at t_1 and $E_2(\lambda; x_i, y_i)$ and $E_2(\lambda; x_i', y_i')$ at t_2, as described in Appendix B.

To provide a point of reference for observed estimates of the mean relative deviation, threshold values were estimated from two previous experiments in which observers discriminated illuminant from non-illuminant changes in simulated samples of Munsell surfaces. Averaged over observers, these estimates varied between about 12% (from Nascimento & Foster, 1997) and 14% (from Nascimento & Foster, 2001), and a criterion level of 13% was used subsequently to define the minimum for detecting violations of ratio invariance.

2.5. Postreceptoral combinations of cone excitations

Four kinds of postreceptoral combinations of cone excitations were considered. Two were varieties of opponent-color transformations producing chromatic and achromatic responses and two were varieties of cone-opponent transformations producing spectrally sharpened responses. Differences in the choice of cone fundamentals proved unimportant.

Let l, m, and s be the cone excitations at an arbitrary point (x, y) in a scene. The first postreceptoral transformation was a theoretical decorrelating operation (Buchsbaum & Gottschalk, 1983). An achromatic response a, a "red–green" chromatic response p, and a "yellow–blue" chromatic response q were obtained by a linear transformation[4]:

$$a = 0.89l + 0.46m,$$
$$p = -0.46l + 0.88m + 0.01s, \tag{2}$$
$$q = -0.01m + 0.99s.$$

The second, nonlinear, postreceptoral transformation was a data-driven decorrelating operation (Ruderman, Cronin, & Chiao, 1998), which, crucially, contained a logarithmic compression:
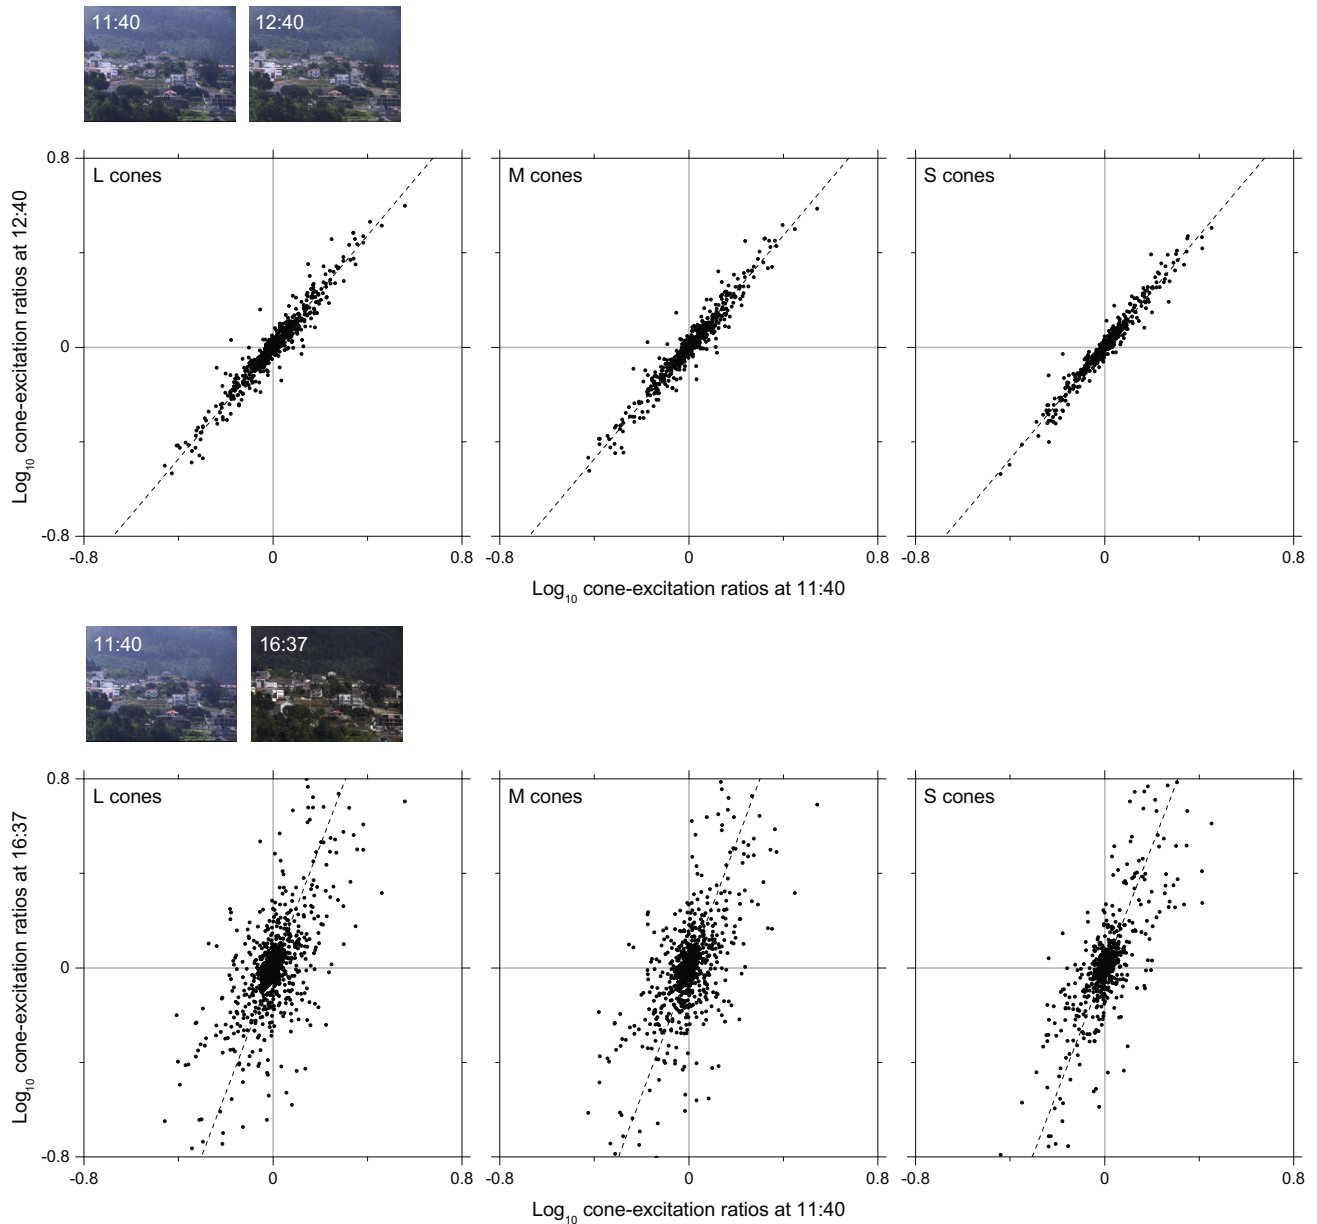
$$a = 0.58 \log(l/\bar{l}) + 0.58 \log(m/\bar{m}) + 0.57 \log(s/\bar{s}),$$
$$p = 0.70 \log(l/\bar{l}) - 0.71 \log(m/\bar{m}) + 0.01 \log(s/\bar{s}), \tag{3}$$
$$q = 0.41 \log(l/\bar{l}) + 0.40 \log(m/\bar{m}) - 0.82 \log(s/\bar{s}),$$

where $\bar{l}$, $\bar{m}$, and $\bar{s}$, are sample means.

The third, linear, postreceptoral transformation maximized spectral sharpening (Foster & Snelgar, 1983) of cone-opponent responses $l^{\#}$, $m^{\#}$, and $s^{\#}$, subject to certain constraints (Finlayson, Drew, & Funt, 1994):

$$l^{\#} = 2.46l - 1.97m + 0.08s,$$
$$m^{\#} = -0.66l + 1.58m - 0.12s, \tag{4}$$
$$s^{\#} = 0.09l - 0.14m + 1.00s.$$

---

[1] To preserve a neutral points-based approach to sampling, images were not segmented into separate surfaces (Foster et al., 2009). Samples within a few pixels of the image boundary were excluded to avoid artifacts from the registration process.

[2] The notation is nonstandard to avoid confusion with that for cone spectral sensitivities.

[3] Spectral radiances E_t were expressed in W m^{-2} sr^{-1} nm^{-1} and spectral sensitivities $L(\lambda)$, $M(\lambda)$, and $S(\lambda)$ were normalized to a maximum of unity on a linear energy scale.

[4] Numerical values rounded to 0.01.

Fig. 3. Scatter plots of spatial ratios of cone excitations across different times of day for the Nogueiró scene. In the top row, data are shown for successive times of 11:40 and 12:40. For each pair of randomly chosen points in the scene, the logarithm of the ratio at one time is plotted against the logarithm of the ratio at another time for L, M, and S cone excitations. For clarity, just 1000 data points from the 10,000 available are plotted. The spacings of points in each sample pair were drawn randomly from the logarithmic scale 1, 2, 4, ..., 256 pixels, and their orientations randomly from the circle. Image sizes were 1344 × 1024 pixels, corresponding approximately to 6.9° × 5.3°. The dashed lines represent orthogonal linear regressions. In the bottom row, corresponding data are shown for more widely separated times of 11:40 and 16:37.

The fourth postreceptoral transformation, also linear, maximized information retrieved from scenes (Foster, Marín-Franch, Amano, & Nascimento, 2009):

$$l^\# = l - 0.97m + 0.09s,$$
$$m^\# = -0.25l + m - 0.17s, \tag{5}$$
$$s^\# = 0.01l - 0.03m + s.$$

Spatial ratios of these transformed cone excitations at points (x, y) and (x', y') and mean relative deviations between the ratios at different times were calculated as in Section 2.4, except for transformation (3) which involved logarithmic compression. Accordingly, with transformation (3), ratios were replaced by differences, that is,

$$\delta a = a - a',$$
$$\delta p = p - p', \tag{6}$$
$$\delta q = q - q'.$$

Differences in these differences across time (i.e. logarithmic ratios of ratios) were then averaged and exponentiated to enable direct comparisons with mean relative deviations from linear postreceptoral transformations.

2.6. Association between postreceptoral differences

A test was also made of a previously reported association (Fine, MacLeod, & Boynton, 2003) between the logarithmic differences δa, δp, and δq in postreceptoral combinations of cone excitations
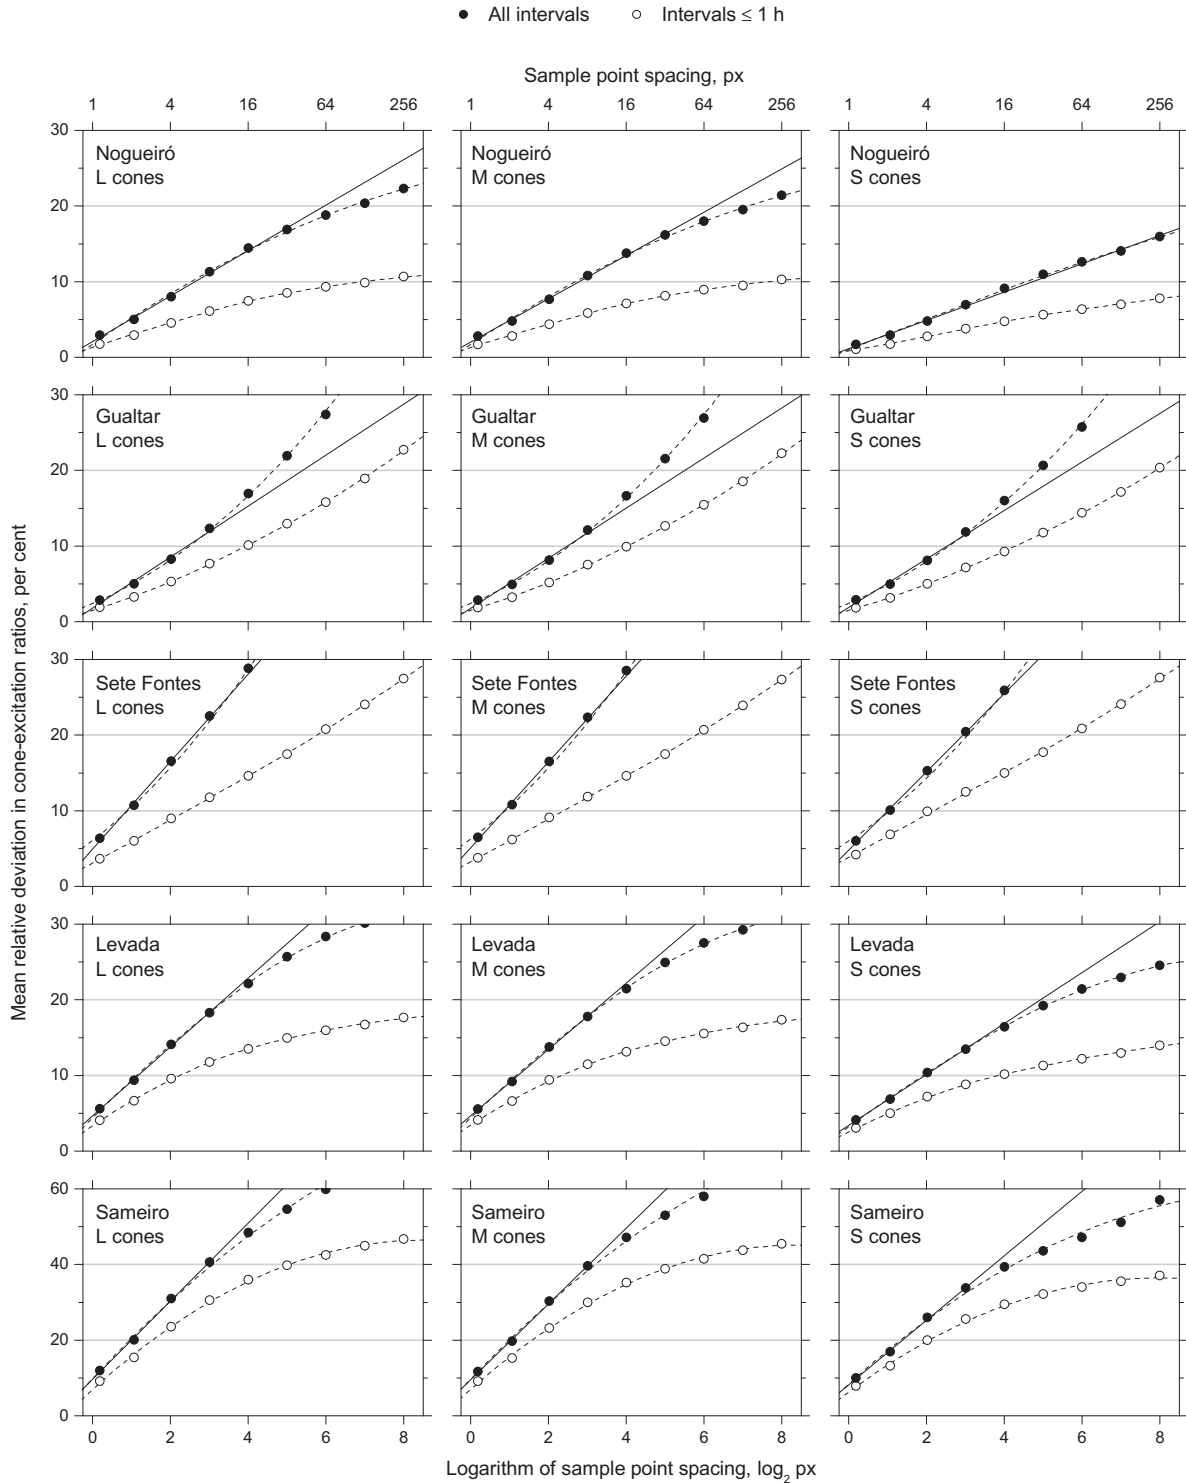
Fig. 4. Deviations in spatial cone-excitation ratios and spacing of points in sample pairs. The symbols represent mean relative deviation in ratios plotted against $\log_2$ spacing in pixels for L, M, and S cone excitations. Image sizes were 1344 × 1024 pixels, corresponding approximately to 6.9° × 5.3°. The solid circles represent means over all (non-zero) time intervals and the open circles means over intervals of up to 1 h. The continuous straight lines are least-squares linear fits over the first four spacings, and the dashed curves are locally weighted quadratic regression fits over all spacings (Cleveland, 1979; Żychaluk & Foster, 2009). The ordinate scale was chosen to show the full or almost the full range of mean relative deviations, but for the Sameiro scene the scale was doubled (see Section 2.3). Estimated SEs of the estimated means were too small to display.

according to definition (6). The pair-wise associations between the differences were quantified by the mutual information I (Cover & Thomas, 2006), that is, $I(\delta a; \delta p)$, $I(\delta p; \delta q)$, and $I(\delta q; \delta a)$. Mutual information was estimated by the asymptotically bias-free k-nearest-neighbor estimator due to Kozachenko and Leonenko (Goria, Leonenko, Mergel, & Novi Inverardi, 2005; Kozachenko & Leonenko, 1987) used in a computationally efficient, offset form (Marín-Franch & Foster, 2013). The correlation between the estimates $I(\delta a; \delta p)$, $I(\delta p; \delta q)$, and $I(\delta q; \delta a)$ over intervals was quantified by the mean Pearson product moment correlation coefficient.

Spacing 8 px

11:44

15:45

Spacing 32 px

11:44

15:45

Fig. 5. Distribution of pairs of sample points in the Gualtar scene under uneven changes in illumination geometry between 11:44 (left) and 15:45 (right). The pairs are indicated by the small red vectors connecting the sample points and were identified according to the method described in Appendix B. The spacings of the sample points were 8 pixels (top) and 32 pixels (bottom), corresponding approximately to 2.4 and 9.7 arcmin angular subtense at the camera.

3. Results and comment

3.1. Spatial ratios of cone excitations

The degree of invariance of ratios of cone excitations across time can be illustrated with scatterplots, as in Fig. 3 for the Nogueiró scene. For each pair of randomly chosen points in the scene, the logarithm of the ratio at one time is shown against the logarithm of the ratio at another time, where the logarithms are to the base 10. The top row is for ratios across successive times of 11:40 and 12:40 and the bottom row for ratios across more widely separated times of 11:40 and 16:37 (see Fig. 2 for the full range of times available). Ratios are shown separately for L, M, and S cones. The dashed lines represent orthogonal linear regressions (major-axis estimation), appropriate for bivariate data (Warton, Wright, Falster, & Westoby, 2006).

For the plots in the top row, the correlation coefficients were 0.96, 0.96, and 0.98 for L, M, and S cones, respectively. More relevantly, the corresponding mean relative deviations in cone-excitation ratios were 5%, 5%, and 3%. These values are almost the same as those obtained with simulated global changes in illumination spectrum, albeit with much greater changes in correlated color temperatures (Nascimento, Ferreira, & Foster, 2002). For the plots in the bottom row, with the longer time interval, the corresponding correlation coefficients were 0.67, 0.67, and 0.76. The corresponding mean relative deviations in cone-excitation ratios were 21%, 21%, and 15%, greater than the criterion level of 13% for detecting violations of ratio invariance (Section 2.4).

In all the plots the slopes of the regressions are greater than unity and in the bottom row the slopes are greater than in the top row, with the greatest slope at short wavelengths (bottom

right). The increase in slopes can be explained by the different compressive effects of haze on the range of ratios available at the different times of day. Details are given in Appendix C.

3.2. Deviations in ratios and sample spacing

The large deviations in some cone-excitation ratios reported in the last section are not peculiar to the Nogueiró scene. The mean relative deviation over all scenes and non-zero time intervals was 24%, a value that increased to 47% when the spacing of points in sample pairs was drawn randomly from a linear scale rather than from the default logarithmic scale (Section 2.4).

How, then, do deviations in ratios depend on sample point spacing? Fig. 4 shows for each scene and cone class, mean relative

Table 1
Threshold sample point spacings (in pixels) averaged over time intervals and cone classes.[a]

	Nogueiró	Gualtar	Sete Fontes	Levada
All time intervals	25.9	10.3	3.2	4.6
Time intervals ⩽1 h[b]	–	40.5	11.3	33.8

[a] Each entry shows the estimated spacing of points in sample pairs yielding a criterion mean relative deviation in cone-excitation ratios of 13%, the minimum for detection (Section 2.4). SEs, not shown, were estimated by bootstrap resampling (Foster & Bischof, 1997) and were each <7% of the corresponding estimated mean. For the Nogueiró scene, an estimated threshold for intervals ⩽1 h was not defined as mean relative deviations were all <13%. Data for the Sameiro scene were omitted for the reasons indicated in Section 2.3. Image sizes were 1344 × 1024 pixels, corresponding approximately to 6.9° × 5.3°.

[b] The 1 h limit was the shortest interval common to all time-lapse sequences (see Fig. 2).
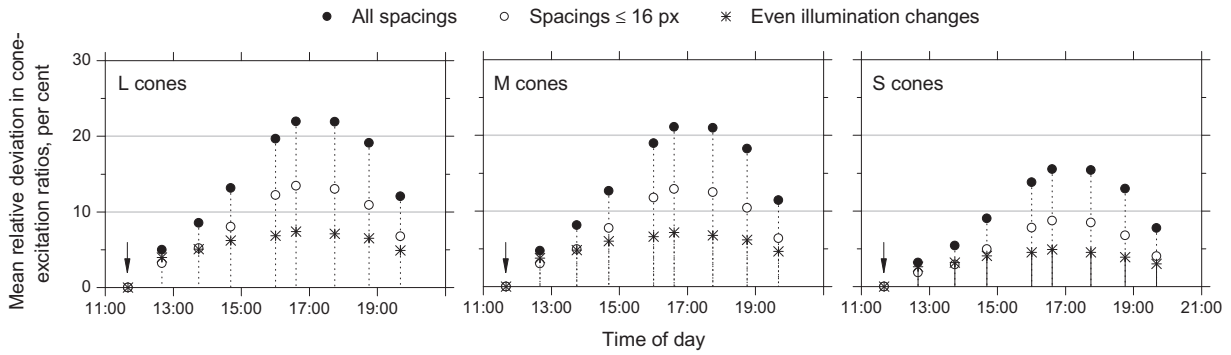
Fig. 6. Deviations between spatial cone-excitation ratios across different times of day for the Nogueiró scene. Each panel shows the mean relative deviation between ratios at 11:40 and over successive, approximately one-hour intervals (the arrow indicating trivial zero deviation at zero interval). The solid circles represent means over all point spacings, i.e. 1, 2, 4, ..., 256 pixels; the open circles means over spacings of up to 16 pixels, corresponding to 4.9 arcmin angular subtense at the camera; and the star symbols means over pairs of points with even changes in illumination geometry (Section 3.2 and Appendix B). Means for L, M, and S cone excitations were based on the same 10,000 pairs of points sampled from the scene of size 1344 × 1024 pixels. Estimated SEs of the means were too small to display.

deviation plotted against the logarithm of the spacing in pixels. The solid circles represent means over all (non-zero) time intervals and the open circles means over intervals of up to 1 h, the shortest interval common to all time-lapse sequences (see Fig. 2). The continuous straight lines are least-squares linear fits over the first four spacings and the dashed curves are locally weighted quadratic regression ('loess') fits over all spacings (Cleveland, 1979; Żychaluk & Foster, 2009). The loess fits, used later to estimate thresholds, accounted for 98% of the variance in the data.

For scenes with vegetation (Nogueiró, Levada, Sameiro) mean relative deviations were somewhat smaller for S cones than for L and M cones, and for each scene the effects of point spacing were similar. But over scenes, the mean relative deviation for a given spacing varied by a factor of two or more. Reasons for the exceptionally large deviations with the scene Sameiro have already been noted (Section 2.3), and data from that scene were not used further.

For each scene and cone class, the dependence of mean relative deviation on log spacing was smooth and monotonic, and almost exactly linear up to log spacings of 3–4 log pixels, i.e. 8–16 pixels. This lawful behavior was obviously not an artifact of the size of the line-spread function of the hyperspectral imaging system (SD approximately 1.3 pixels; see Section 2.1). A monotonic increase with spacing might be expected, since the probability of sampling from an uneven change in illumination geometry should increase with point spacing. A strictly linear increase with log spacing, however, is contingent on specific scene properties, including the relative frequency of different-sized surfaces in each scene (Yang & Purves, 2003).

The relationship between sample point spacing and uneven changes in illumination geometry is illustrated in Fig. 5. It shows color images rendered from two hyperspectral images of the Gualtar scene at 11:44 (left) and 15:45 (right), as in Fig. 2. Sample pairs with uneven changes were identified according to the method described in Appendix B and are denoted in the figure by the small red vectors of length either 8 pixels (top) or 32 pixels (bottom). Uneven changes were clearly more frequent with the larger of the two spacings.

From Fig. 4, it is evident that with sufficiently small spacings of sample points, mean relative deviations fell below the criterion level of 13% for detection (Section 2.4). Threshold values of those spacings were estimated from the loess fits for each scene and cone class. Table 1 shows the mean estimates with and without upper limits on time intervals. There were marked variations over scenes, with larger values for time intervals of up to 1 h. For the Nogueiró scene, no threshold could be defined.

For a more detailed analysis of the effects of time interval, it proved useful to define a common upper limit on point spacings that, over all scenes and time intervals, yielded sample mean relative deviations below the criterion level. An upper limit of 32 pixels yielded 13.3% and one of 16 pixels yielded 11.3%. Neither of these limits has in itself any absolute significance, for as Table 1 showed, threshold spacings depend strongly on the scenes themselves. A common upper limit of 16 pixels was adopted.
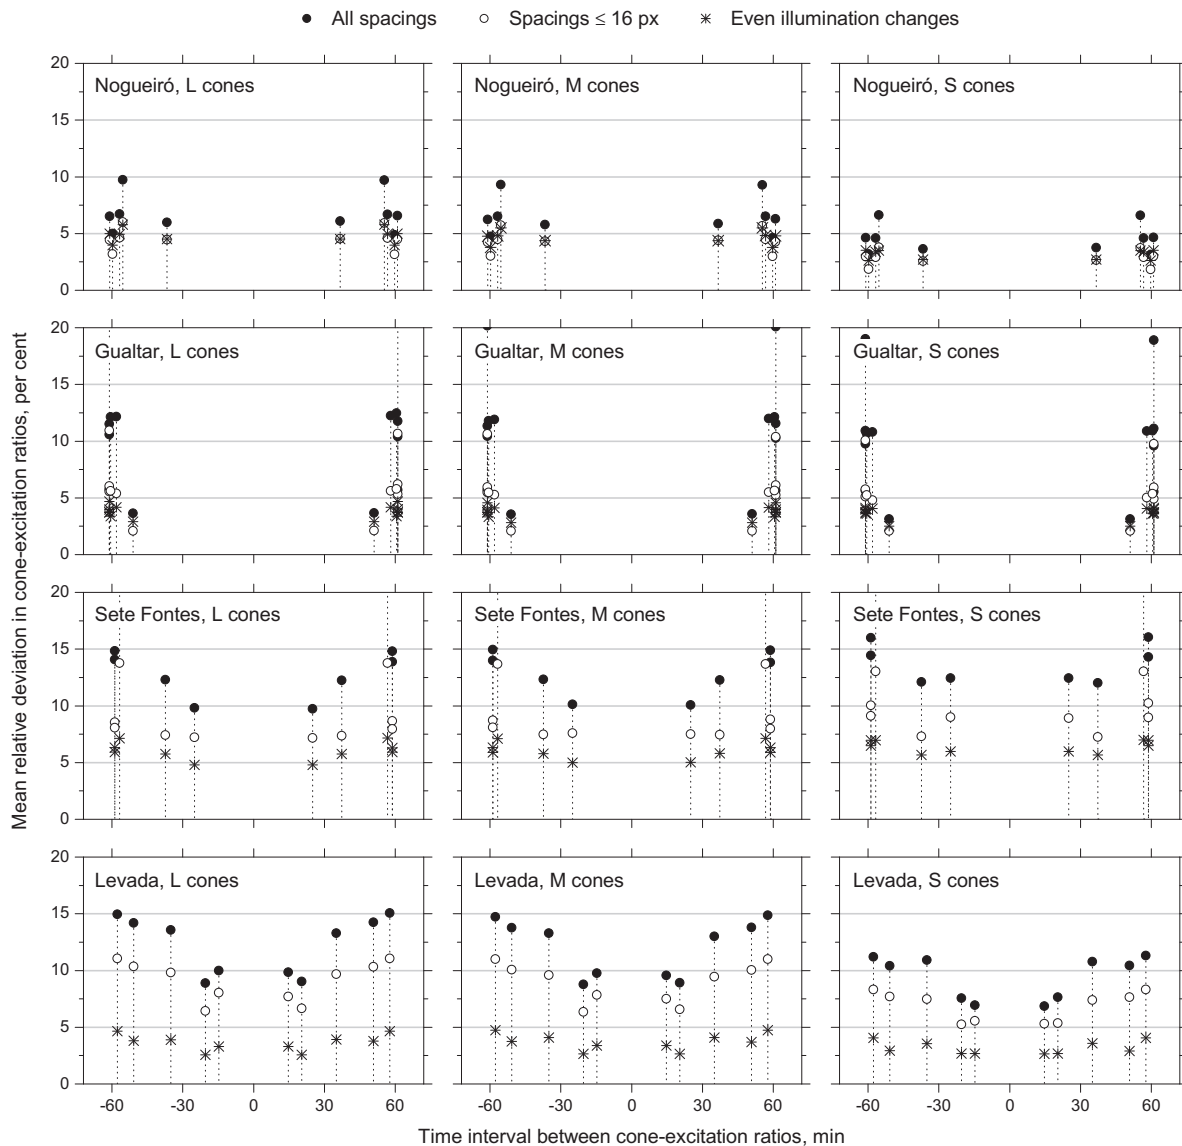
3.3. Deviations in ratios across time

How do differences between cone-excitation ratios vary according to the time of day? The plots in Fig. 6 show for the Nogueiró scene the mean relative deviation between ratios at an initial reference time of 11:40 and at successive, approximately one-hour time intervals.

The solid circles represent means over all spacing of points in sample pairs, the open circles over spacings of up to 16 pixels (the common upper limit, Section 3.2), and the star symbols over pairs of points with even changes in illumination geometry (Appendix B). Tellingly, over intervals of up to 2 h, means for spacings of up to 16 pixels were similar in magnitude to those for points with even changes in illumination geometry.

Means over all point spacings initially increased with increasing time interval, although towards the end of the day, the differences in the illuminations decreased, leading to smaller deviations in ratios. Broadly similar dependencies on time interval, not shown here, were found with an initial reference time of 12:40, and with the next one, and so on; and likewise for the Gualtar, Sete Fontes, and Levada scenes.

Over all reference times and scenes, there were 33 plots of the form shown in Fig. 6, but the data most relevant to the invariance of ratios are from the shorter time intervals around each reference time. Fig. 7 shows mean relative deviation plotted against signed time interval for all non-zero intervals of up to 1 h. As in Fig. 6, the solid circles represent means over all point spacings, the open circles over spacings of up to 16 pixels, and the star symbols over pairs of points with even changes in illumination geometry. The additional data points for the Setes Fontes and Levada scenes at intervals of less than 30 min were derived from acquisitions at the end of the day when the light level was changing rapidly (Section 2.1). Data values were not exactly symmetric about zero time interval because each was from an independent set of sample points.

For sufficiently short time intervals, mean relative deviations did not exceed the criterion level of 13% for detection

Fig. 7. Deviations between spatial cone-excitation ratios across shorter time intervals. Each panel shows the mean relative deviation in ratios plotted against the signed time interval between ratios for intervals of up to 1 h. The number of data points depends on the imaging conditions for each scene. Data for zero time intervals have been omitted. For each scene, the ordinate scale has been adjusted to show the full or almost the full range of mean relative deviations (extreme values at the longer time intervals have been clipped). Other details as for Fig. 6.

Table 2

Mean relative deviations (%) in spatial ratios of excitations of long-, medium-, and short-wavelength-sensitive (L, M, and S) cones with and without upper limits on point spacings and time intervals.[a]

		L	M	S
All intervals	All spacings	21.2	20.8	17.7
	Spacings ⩽16 pixels[b]	12.1	11.8	9.9
Intervals ⩽1 h[c]	All spacings	11.3	11.1	9.8
	Spacings ⩽16 pixels[b]	7.1	7.0	6.2

[a] Each entry shows the estimated mean relative deviation taken over scenes, spacings of sample points, and time intervals. The mean for each individual scene was estimated from 10,000 pairs of scene points drawn randomly from the scene for each spacing and time interval, and then averaged over all spacings and (non-zero) time intervals. SEs, not shown, were each <2% of the corresponding estimated mean.

[b] The 16 pixels limit on spacings was derived from the average estimated threshold spacing for detecting violations of ratio invariance (Section 3.2). Image sizes were 1344 × 1024 pixels, corresponding approximately to 6.9° × 5.3°.

[c] The 1 h limit was the shortest interval common to all time-lapse sequences (see Fig. 2).

Table 3

Mean relative deviations (%) in extreme spatial ratios of excitations of long-, medium-, and short-wavelength-sensitive (L, M, and S) cones for all illumination changes and for illumination changes with even geometry.[a]

		L	M	S
All illumination changes	All intervals	62.6	62.1	59.2
	Intervals ⩽1 h	32.6	32.6	32.5
Even illumination changes	All intervals	8.7	8.7	9.2
	Intervals ⩽1 h	8.3	8.4	8.8

[a] Ratios were either <0.32 or >3.2 . Illumination changes with even geometry are defined in Appendix B. Other details as for Table 2.

(Section 2.4), whether or not point spacings were limited to 16 pixels. Table 2 shows the means over scenes with and without upper limits on point spacings and time intervals.[5]

_____

[5] Slightly smaller mean relative deviations were obtained with a 4 pixels upper limit on spacings of sample points, as reported at VSS 2015.

Table 4
Mean relative deviations (%) in spatial ratios of excitations of long-, medium-, and short-wavelength-sensitive (L, M, and S) cones for illumination changes with even geometry.[a]

		L	M	S
All intervals	All spacings	5.8	5.7	5
	Spacings ⩽16 pixels	4.7	4.7	3.9
Intervals ⩽1 h	All spacings	4.6	4.5	4.1
	Spacings ⩽16 pixels	3.5	3.4	3.1

[a] Illumination changes with even geometry are defined in Appendix B. Other details as for Table 2.

3.4. Extreme ratios

Is the invariance of spatial cone-excitation ratios, where it occurs, an artifact of most ratios being close to unity? As the scatterplots in the top row of Fig. 3 illustrate for the Nogueiró scene, the distributions of the logarithms of the ratios are concentrated around zero, and it might be hypothesized that such values are likely to arise from points sampled from single surfaces under similar illumination conditions. Ratios from those points are therefore likely to remain constant.

As the scatterplots demonstrate, some extreme ratios are preserved (those represented by data points in the top right and bottom left of the plots). But to test this hypothesis properly, mean relative deviations were calculated over all scenes, time intervals, and sample point spacings for which log ratios fell outside the range [−0.5, 0.5], i.e. for which ratios were either less than 0.32 or greater than 3.2. For each scene, these extreme ratios were in the minority, constituting about 5% of the total, and the number of sample pairs was accordingly increased to 100,000.

Table 3 shows the means over scenes for extreme ratios from all illumination changes and from illumination changes with even geometry (Appendix B). Providing that uneven changes were excluded, the means fell well below the criterion level of 13% for detecting violations of ratio invariance (Section 2.4).

3.5. Postreceptoral ratios

For the three linear postreceptoral transformations of cone excitations described in Section 2.5, mean relative deviations were closely similar to those shown in Table 2 for untransformed cone excitations. Tables D.1–D.4 in Appendix D show the means for both linear and nonlinear transformations.

For the linear opponent-color transformation of definition (2), the linear spectral sharpening transformation of definition (4), and the linear transformation of definition (5), which maximized information retrieved, differences in mean relative deviations with respect to those for the untransformed cone excitations (Table 2) were less than 2 percentage points for each sampling condition. But for the combination of logarithmic compression and linear transformation of definition (3), mean relative deviations were much higher for achromatic responses, exceeding the criterion level of 13% for detection, and, conversely, much smaller for the two chromatic responses (Appendix D, Table D.2).

The pairwise associations between the logarithmic differences δa, δp, and δq in postreceptoral combinations defined by the mutual information quantities $I(\delta a; \delta p)$, $I(\delta p; \delta q)$, and $I(\delta q; \delta a)$ (Section 2.6) were correlated across time. Pearson product moment correlation coefficients averaged across time intervals and spacings of sample points ranged from 0.45 for the Levada scene to 0.89 for the Gualtar scene.

3.6. Natural and simulated illumination changes

Restricting pairs of sample points to those with even changes in illumination geometry (Appendix B) provided a rough test of previous simulations with global changes in illumination spectrum and no accompanying change in illumination geometry (e.g. Nascimento, Ferreira, & Foster, 2002). Table 4 shows mean relative deviations from pairs of points with even changes in illumination geometry. Depending on the sampling constraints, the means ranged from 3.1% to 5.8%, overlapping the values of 3.1–4.8% obtained with simulated global changes in illumination spectrum (Nascimento, Ferreira, & Foster, 2002, Table 1). There are, nonetheless, several additional factors that need to be considered in this comparison, as explained in Section 4.

4. Discussion

Because of the complexity of natural scenes, geometrical variations in local illumination can be similar in size to the variations in global illumination over the course of the day (Nascimento, Amano, & Foster, 2016). Whether signals such as the spatial ratio of cone excitations generated from reflected light are invariant across time depends on the interaction of these spatial and temporal factors, which varies from scene to scene.

Evaluated over scenes, time intervals, and spacings of sample points, the estimated mean relative deviation in ratios was about 24%, a value that increased to 47% when sample spacings were drawn from a linear rather than logarithmic scale. These estimates are greater than the criterion level of 13% for detecting violations of ratio invariance and much greater than the level of about 4% found with simulated global changes in illumination spectrum (Nascimento, Ferreira, & Foster, 2002). Yet when sampling was constrained, either to moderate time intervals or to short distances, the estimated mean relative deviation decreased to about 11%, below the criterion level for detection.

Unsurprisingly, the spacing of sample points at which deviations in ratios reached the criterion level varied markedly over scenes. Estimates ranged from 3 to 26 pixels (1–8 arcmin), when averaged over all time intervals, and from 11 pixels (3 arcmin) to unlimited values when averaged over intervals of up to 1 h. While these threshold spacings might be gauged intuitively in relation to the overall image size of 1344 × 1024 pixels (corresponding approximately to 6.9° × 5.3°), they depend, as implied earlier, more on the physical distances within the scenes themselves than on the angle they subtended at the camera.

It should not be concluded, however, that in the absence of spatial and temporal sampling constraints, deviations in ratios are caused by changes in illumination geometry, per se. Rather, they are caused by the unevenness in those changes where that unevenness exists. For the Nogueiró, Gualtar, Sete Fontes, and Levada scenes, the proportion of samples from uneven illumination changes ranged from 10% to 66%. The remaining samples were from even illumination changes, but still involved movement of the sun and possibly cloud and variations in atmospheric scatter, mutual illumination, and local attached and cast shadows. Revealingly, estimates of mean relative deviations in ratios from even changes fell to levels found with simulated global changes in illumination spectrum.

Nevertheless the similarity in the two kinds of estimates, that is, from natural illumination changes with even geometry and from simulated global changes in illumination spectrum, should be interpreted with care. The simulations (Nascimento, Ferreira, & Foster, 2002) were based on larger ranges of daylight illuminants, many more natural scenes, and linear sampling of point spacings. No allowance was made for independent sampling noise (e.g.

Osorio & Vorobyev, 2005; Vorobyev & Osorio, 1998). Furthermore, the global changes effectively averaged illumination shifts that at the end of the day could have been in opposite directions in directly illuminated and shadowed regions of the scene (Hubel, 2000), although control simulations in which shadowed regions were masked suggested that averaging did not distort estimates (Foster et al., 2006, Sect. 3.F). Overall, these simulations appear not to seriously misrepresent the effects of natural variations in illumination, providing that uneven changes in illumination geometry are excluded.

With natural scenes, one of the factors affecting the invariance of ratios from distant regions is atmospheric scatter, which could have blurred the changing distribution of shadows. Thus haze may have contributed to the lower deviations in ratios from the Nogueiró scene early and later in the day (Fig. 2, top left, 11:40, 12:40, and 13:45), even though it covered only part of the scene. Yet it was largely absent at other times and the low mean relative deviations in ratios from that scene were probably due more to the properties of the scene itself. Haze is visible over the distant hills in the Levada scene, but it, too, seems to have had a limited impact on deviations in ratios from the scene as a whole.

It was anticipated that ratios of linear combinations of cone excitations, both opponent-color and cone-opponent, would yield levels of invariance similar to those found with ratios of cone excitations. A previous analysis of ratios of opponent-color signals revealed an analogous result with simulated global changes in illumination spectrum (Nascimento & Foster, 2000). Interestingly, an association previously reported between differences in logarithmic combinations of cone excitations (Fine, MacLeod, & Boynton, 2003) was preserved over time intervals, most strongly for the Gualtar scene and least so for the Levada scene.

As already noted, the approximate invariance of ratios under global changes in illumination spectrum has been invoked in the explanation of a broad range of perceptual color phenomena, not least the phenomenon of color constancy in its various forms, including relational and displacement color constancies (see, in addition to previously cited sources, Amano & Foster, 2004; Chittka, Faruq, Skorupski, & Werner, 2014; Foster et al., 2001; Hurlbert & Wolf, 2004; Kamermans, Kraaij, & Spekreijse, 1998; Khang & Zaidi, 2002; Kulikowski et al., 2012; Lee, Dawson, & Smithson, 2012; Plet & Gerbino, 2001; Ripamonti & Westland, 2003; Tokunaga & Logvinenko, 2010; Wachtler, Albright, & Sejnowski, 2001; Werner, 2014; Westland & Ripamonti, 2000). Moreover, a generalization of spatial ratios that replaces scaling by an affine transformation has been exploited in modeling the additional effects of changes in viewing medium (Hagedorn & D'Zmura, 2000; Romero, Luzón-González, Nieves, & Hernández-Andrés, 2011; see also Zaidi, 1998). The finding of an approximate invariance of ratios under constrained sampling may extend the scope of these and related studies to environments undergoing natural changes in illumination. It is, though, important to distinguish between optical limits on the invariance of cone-excitation ratios or their equivalents and limits on performance set more centrally (e.g. Arnold & Chittka, 2012; Cornelissen & Brenner, 1995; Kentridge, Heywood, & Cowey, 2004; Linnell & Foster, 1996; Murray et al., 2014; Rinner & Gegenfurtner, 2002; Smithson & Zaidi, 2004; Wachtler, Albright, & Sejnowski, 2001; Werner, 2014). The present work has been concerned solely with optical limits. It is also important to distinguish between the information extracted at points sampled randomly from scenes and the information extracted from particular areas, such as object edges (Gijsenij, Gevers, & van de Weijer, 2012; Peyvandi, Nieves, & Gilchrist, 2013).

As a final cautionary note, these estimates of mean relative deviations in cone-excitation ratios may represent only an upper bound on the true invariance of ratios in natural scenes. During image acquisition both movement within scenes and movement of the camera itself could have introduced errors. Large scene movements were easily detected and the images then rejected, as with the Sameiro scene, and small camera movements were automatically compensated for by image processing, at least up to the resolution of the registration routines. Even so, small undetected movements within scenes and other imaging artifacts would have inflated any differences in observed ratios.

These caveats aside, it is evident that within certain sampling constraints, spatial ratios of cone excitations, and also of opponent-color combinations, provide an approximately invariant signal for stable surface-color inferences, despite a transient optical environment where changes in the spectrum of the illumination are accompanied by changes in illumination geometry.

The hyperspectral radiance images used in this study are available from the authors' web sites, http://personalpages.manchester.ac.uk/staff/d.h.foster/ and http://online.uminho.pt/pessoas/smcn/.

Acknowledgments

Appendix A. Quantifying deviations in ratios

With the notation of Section 2.4, the mean relative deviation MRD between ratios $r_{i,1}$ and $r_{i,2}$ at times t_1 and t_2, respectively, taken over sample pairs $i = 1, 2, \ldots, N$ may be formulated in two ways:

$$MRD_A = \frac{1}{N} \sum_{i=1}^{N} \frac{|r_{i,1} - r_{i,2}|}{\min\{r_{i,1}, r_{i,2}\}} \qquad (A.1)$$

and

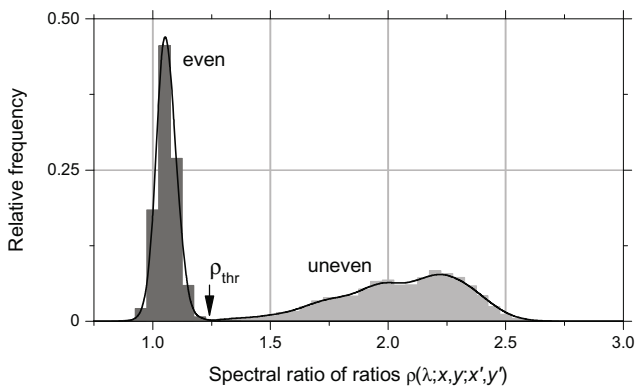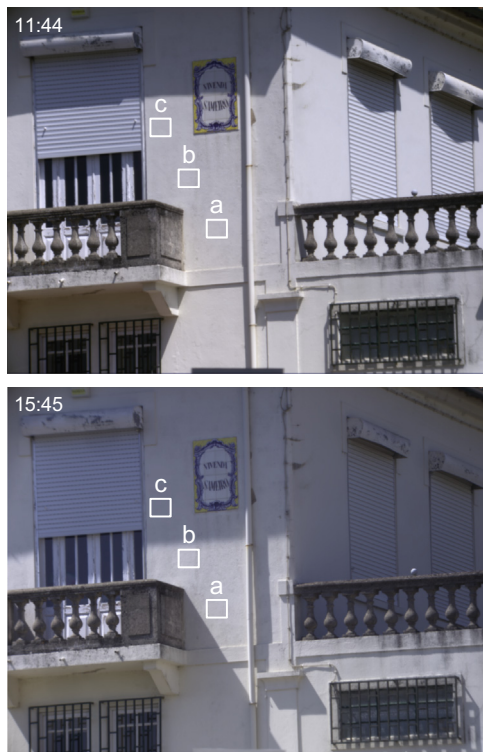$$MRD_B = \frac{1}{N} \sum_{i=1}^{N} \frac{|r_{i,1} - r_{i,2}|}{(r_{i,1} + r_{i,2})/2}. \qquad (A.2)$$

The formulation MRD_A of definition (A.1) has been used previously with global changes in illumination spectrum (e.g. Nascimento, Ferreira, & Foster, 2002). Unfortunately, with additional changes in illumination geometry, the estimate provided by MRD_A can become biased when it contains terms whose denominators are much smaller than the numerators, for example, when just one member of the pairs (x_i, y_i) and (x'_i, y'_i) is in shadow at t_1 and not at t_2. The bias can be reduced by replacing the minimum operator in the denominator by the mean, as in the formulation MRD_B of definition (A.2). When multiplied by 100 it is the symmetric mean absolute percentage error.

By construction, MRD_A is never smaller than MRD_B. In fact, $MRD_A/MRD_B = 1 + MRD_A/2$. Yet in simulations with natural scenes undergoing global changes in illumination spectrum, the two measures were found to differ little from each other. For example, if the illuminant changed from a daylight with correlated color temperature of 25,000 K to one of 4300 K, the value of MRD_A was, at most, 4.8% (Nascimento, Ferreira, & Foster, 2002, Table 1). The corresponding value of MRD_B is then 4.7%. But the difference between the two definitions becomes significant when one of the ratios is very different from the other. If $r_{i,1}$ has some particular nonzero value and $r_{i,2}$ tends to zero, then the summand in definition (A.1) tends to infinity, whereas the summand in definition (A.2) tends to 2, with little impact on MRD_B. In this report, all values of the mean relative deviation were calculated according to definition (A.2), shown as definition (1) in the main text.

Appendix B. Classifying changes in illumination geometry

There are many ways of identifying shadowed regions in image data (see e.g. Finlayson, Hordley, Lu, & Drew, 2006; Gijsenij et al., 2012; Gu & Robles-Kelly, 2014; Jiang, Schofield, & Wyatt, 2011; Párraga, Brelstaff, Troscianko, & Moorehead, 1998; Salvador, Cavallaro, & Ebrahimi, 2004; Tappen, Freeman, & Adelson, 2005'). The requirement in this analysis, however, was more specific, namely to identify pairs of sample points with uneven changes in illumination geometry. These changes may be a result of a change in the distribution of shadows, but could also occur with strong differences in mutual illumination.

The Gualtar building scene provides a simple example. The color images in Fig. B.1 were rendered from hyperspectral images acquired at 11:44 and 15:45. Rectangular sample areas a, b, and c have been outlined in white. The change in illumination from 11:44 to 15:45 across the two areas a and b was even: spatial ratios of cone excitations from a and b depend only on the spectrum of the illumination and on the surface spectral reflectances, for the given viewing direction and angular distribution of incident illumination (more precisely on the spectral bidirectional reflectance distribution function, BRDF; Nicodemus, Richmond, Hsia, Ginsberg, & Limperis, 1997). In contrast, the change in illumination from 11:44 to 15:45 over the areas a and c (and b and c) was uneven: ratios of cone excitations from a and c at 11:44 are unaffected by shadow whereas those at 15:45 are affected strongly. The ratios from a and c can therefore differ arbitrarily across different times. This uncertainty is of course not restricted to planar surfaces of buildings and cast shadows; it also occurs with natural non-planar materials and attached shadows (see Fig. 2 for examples).

To provide an objective criterion for distinguishing between even and uneven changes in illumination geometry at two arbitrary points in a scene, a test statistic was constructed from the reflected spectral radiances at those points. With the notation of Section 2.4, let $E_1(\lambda; x, y)$ and $E_1(\lambda; x', y')$ be the respective spectral radiances at (x, y) and (x', y') at time t_1 and $E_2(\lambda; x, y)$ and $E_2(\lambda; x', y')$ the corresponding values at time t_2. The ratios $E_1(\lambda; x, y)/E_1(\lambda; x', y')$ and $E_2(\lambda; x, y)/E_2(\lambda; x', y')$ each depend on the spectral reflectances of the two surfaces, but if the illumination change is even, as over areas a and b in Fig. B.1, then the ratio of these ratios:

$$\rho(\lambda; x, y; x', y') = \frac{E_1(\lambda; x, y)}{E_1(\lambda; x', y')} \cdot \frac{E_2(\lambda; x', y')}{E_2(\lambda; x, y)} \tag{B.1}$$

should be close to unity at all wavelengths λ, with any small departures depending on the details of the spectral BRDF (this ratio of ratios of spectra should not be confused with the ratio of ratios of cone excitations; see e.g. Brill & West, 1981; van Trigt, 2005). If, though, the illumination change is uneven, as over areas a and c, then $\rho(\lambda; x, y; x', y')$ should depart sharply from unity over a range of wavelengths λ. The plot in Fig. B.1 shows the empirical distributions of $\rho(\lambda; x, y; x', y')$ for the paired areas a, b, and a, c. The intersection of the two distributions, at $\rho_{thr} = 1.24$, provides a criterion for deciding whether the sample for a particular pair of points elsewhere in the scene is from an even change or not.

For each of the four other scenes used in this study, empirical distributions of $\rho(\lambda; x, y; x', y')$ were obtained by selecting sample areas with even and uneven illumination changes similar to those in Fig. B.1 for the Gualtar scene. The areas varied from side 10 pixels to 50 pixels. Values of ρ_{thr} ranged from 1.22 to 1.66 over the five scenes.

A separate estimate of the empirical distribution of $\rho(\lambda; x, y; x', y')$ was obtained from three scenes where there was a sufficiently large embedded neutral sphere (see e.g. bottom left of the Levada scene, Fig. 2). On each sphere, sample areas similar to those in Fig. B1 were identified, but their position and size were allowed to vary. Values of ρ_{thr} ranged from 1.31 to 1.40. For consistency over scenes, a common value of ρ_{thr} was used and set equal to the mean of 1.34. With this criterion, the proportion of sample pairs with even illumination changes varied over the four scenes from 34% for the Levada scene to 90% for the Gualtar scene. The proportion was 1% for the fifth Sameiro scene.

To provide an independent test of whether this method of classification could have falsely excluded pure spectral changes in illumination, it was applied to hyperspectral images of the same scenes undergoing simulated global changes in illumination spectrum, as in Nascimento et al. (2002). A hyperspectral image was taken from each time-lapse sequence from each scene and its radiances $E(\lambda; x, y)$ were converted to effective spectral reflectances $R(\lambda; x, y)$ by scaling with a known reflectance $R_0(\lambda; x_0, y_0)$ at some point (x_0, y_0) in the scene (see e.g. Foster et al., 2006, Appendix A), thus,

Fig. B.1. Classifying even and uneven changes in illumination geometry. The two color images were rendered from hyperspectral radiance images of the Gualtar scene at 11:44 and 15:45. The plot below shows empirical distributions of the statistic $\rho(\lambda; x, y; x', y')$ given by definition (B.1) for even (left) and uneven (right) illumination changes across scene points drawn respectively from areas a and b and from a and c, as indicated in the color images.

$$R(\lambda; x, y) = E(\lambda; x, y) \frac{R_0(\lambda; x_0, y_0)}{E(\lambda; x_0, y_0)}.$$

Global daylight illumination changes were simulated by multiplying $R(\lambda; x, y)$ by daylight illuminants with correlated color temperatures of 25,000 K and 4000 K generated from a principal components analysis (Judd, MacAdam, & Wyszecki, 1964). Neither this method of generation nor the choice of source data was critical. Deviations in cone-excitation ratios across the two simulated radiance images were calculated with and without uneven illuminant changes. No difference was detected with any of the five scenes. Since the simulated changes in illumination spectra were greater than any of the natural changes in illumination spectra in the time-lapse sequences, it seems likely that the test statistic $\rho(\lambda; x, y; x', y')$ excluded only sample points where there was also an uneven change in illumination geometry, though not necessarily all such points.

Appendix C. Haze and orthogonal regressions on ratios

The scatterplots of spatial ratios of cone excitations in Fig. 3 have orthogonal regression slopes greater than unity, a property that may be explained by the compressive effects of haze on the range of ratios available. Assume without loss in generality that the ratios are of L-cone excitations. The characteristic bright veil over distant parts of the scene, also known as airlight, is produced by light from the sun and sky not otherwise imaged by the camera being scattered into the imaging path. This light can be represented as a constant additive term, $k > 0$, which depends on both wavelength and viewing distance (Middleton, 1960). Suppose that in the absence of haze the (positive) cone excitations at arbitrary scene points (x, y) and (x', y') are l and l', respectively, and that in the presence of haze they are $l + k$ and $l' + k$, respectively. Suppose that the ratio $l/l' > 1$. Then the ratio $(l+k)/(l'+k)$ satisfies the conditions that $(l+k)/(l'+k) > 1$ and that $(l+k)/(l'+k) < l/l'$. Conversely, suppose that the ratio $l/l' < 1$. Then the ratio $(l+k)/(l'+k)$ satisfies the conditions that $(l+k)/(l'+k) < 1$ and that $(l+k)/(l'+k) > l/l'$. The transformation of l/l' to $(l+k)/(l'+k)$ is therefore strictly monotonic decreasing, i.e. compressive. Let $v = \log(l/l')$ and $u = \log[(l+k)/(l'+k)]$. Since taking logarithms preserves monotonicity, the transformation of v to u is also compressive. For a sample of $2N$ scene points (x_i, y_i) and (x'_i, y'_i), with $i = 1, 2, \ldots, N$, the slope of the orthogonal regression of v_i on u_i is greater than unity, consistent with the plots in Fig. 3 showing slopes in the bottom row greater than in the top row and with the greatest slope for S-cone excitations (Burton & Moorhead, 1987).

Appendix D. Deviations in postreceptoral ratios

Tables D.1–D.4 show mean relative deviations in spatial ratios of combinations of cone excitations for four kinds of postreceptoral transformations (Section 2.5).

Table D.1
Mean relative deviations (%) in spatial ratios of opponent-color combinations A, P, and Q giving achromatic and chromatic red–green and yellow–blue responses, respectively, according to Buchsbaum and Gottschalk (1983).[a]

		A	P	Q
All intervals	All spacings	21.1	20.2	17.7
	Spacings ⩽16 pixels	12.0	11.6	9.9
Intervals ⩽1 h	All spacings	11.2	11.0	9.8
	Spacings ⩽16 pixels	7.0	7.0	6.2

[a] Opponent-color transformations as in definition (2), main text. Other details as for Table 2.

Table D.2
Mean relative deviations (%) in spatial ratios of logarithmic opponent-color combinations A, P, and Q giving achromatic and chromatic red–green and yellow–blue responses, respectively, according to Ruderman, Cronin, and Chiao (1998).[a]

		A	P	Q
All intervals	All spacings	35.2	0.8	5.8
	Spacings ⩽16 pixels	19.7	0.5	3.9
Intervals ⩽1 h	All spacings	18.7	0.5	3.7
	Spacings ⩽16 pixels	11.6	0.4	2.8

[a] Opponent-color transformations as in definition (3), main text. Other details as for Table 2.

Table D.3
Mean relative deviations (%) in spatial ratios of cone-opponent combinations $L^\#$, $M^\#$, and $S^\#$ of cone excitations giving maximally sharpened spectral responses, according to Finlayson, Drew, and Funt (1994).[a]

		$L^\#$	$M^\#$	$S^\#$
All intervals	All spacings	22.0	20.6	17.6
	Spacings ⩽16 pixels	12.5	11.8	9.8
Intervals ⩽1 h	All spacings	11.6	11.0	9.8
	Spacings ⩽16 pixels	7.3	6.8	6.1

[a] Cone-opponent transformations as in definition (4), main text. Other details as for Table 2.

Table D.4
Mean relative deviations (%) in spatial ratios of cone-opponent combinations $L^\#$, $M^\#$, and $S^\#$ of cone excitations maximizing information retrieved, according to Foster et al. (2009).[a]

		$L^\#$	$M^\#$	$S^\#$
All intervals	All spacings	22.8	21.6	17.6
	Spacings ⩽16 pixels	13.0	12.4	9.8
Intervals ⩽1 h	All spacings	12.4	11.5	9.8
	Spacings ⩽16 pixels	7.9	7.3	6.1

[a] Cone-opponent transformations as in definition (5), main text. Other details as for Table 2.

References

Amano, K., & Foster, D. H. (2004). Colour constancy under simultaneous changes in surface position and illuminant. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 271,* 2319–2326.

Amano, K., Foster, D. H., & Nascimento, S. M. C. (2005). Minimalist surface-colour matching. *Perception, 34,* 1009–1013.

Arend, L. (2001). Environmental challenges to color constancy. In B. E. Rogowitz & T. N. Pappas (Eds.). *Human Vision and Electronic Imaging VI, Proceedings of SPIE* (Vol. 4299, pp. 392–399). Bellingham WA: SPIE.

Arnold, S. E. J., & Chittka, L. (2012). Illumination preference, illumination constancy and colour discrimination by bumblebees in an environment with patchy light. *Journal of Experimental Biology, 215,* 2173–2180.

Bloj, M. G., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature, 402,* 877–879.

Brill, M., & West, G. (1981). Contributions to the theory of invariance of color under the condition of varying illumination. *Journal of Mathematical Biology, 11,* 337–350.

Brill, M. H. (2008). Minimal Von Kries illuminant invariance. *Color Research and Application, 33,* 320–323.

Buchsbaum, G., & Gottschalk, A. (1983). Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 220,* 89–113.

Burton, G. J., & Moorhead, I. R. (1987). Color and spatial structure in natural scenes. *Applied Optics, 26,* 157–170.

Carter, G. A., & Knapp, A. K. (2001). Leaf optical properties in higher plants: Linking spectral characteristics to stress and chlorophyll concentration. *American Journal of Botany, 88,* 677–684.

Castiello, U. (2001). Implicit processing of shadows. *Vision Research, 41,* 2305–2309.

Cavanagh, P., & Leclerc, Y. G. (1989). Shape from shadows. *Journal of Experimental Psychology-Human Perception and Performance, 15,* 3–27.

Chittka, L., Faruq, S., Skorupski, P., & Werner, A. (2014). Colour constancy in insects. *Journal of Comparative Physiology A, 200,* 435–448.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association, 74*, 829–836.

Cornelissen, F. W., & Brenner, E. (1995). Simultaneous colour constancy revisited: An analysis of viewing strategies. *Vision Research, 35*, 2431–2448.

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory*. Hoboken, New Jersey: John Wiley & Sons Inc.

Dörrer, D., & Newton, C. (2007). *Unpublished MPhys Project Report*. School of Physics and Astronomy: University of Manchester, UK.

Ekpenyong, N. (2013). *Hyperspectral imaging: Calibration and applications with natural scenes* (Ph.D. thesis). Manchester, UK: University of Manchester.

Elder, J. H., Trithart, S., Pintilie, G., & MacLean, D. (2004). Rapid processing of cast and attached shadows. *Perception, 33*, 1319–1338.

Endler, J. A. (1993). The color of light in forests and its implications. *Ecological Monographs, 63*, 1–27.

Fine, I., MacLeod, D. I. A., & Boynton, G. M. (2003). Surface segmentation based on the luminance and color statistics of natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision, 20*, 1283–1291.

Finlayson, G. D., Drew, M. S., & Funt, B. V. (1994). Spectral sharpening: sensor transformations for improved color constancy. *Journal of the Optical Society of America A – Optics Image Science and Vision, 11*, 1553–1563.

Finlayson, G. D., Hordley, S. D., Lu, C., & Drew, M. S. (2006). On the removal of shadows from images. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 28*, 59–68.

Foster, D. H., Amano, K., Nascimento, S. M. C., & Foster, M. J. (2006). Frequency of metamerism in natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision, 23*, 2359–2372.

Foster, D. H., & Bischof, W. F. (1997). Bootstrap estimates of the statistical accuracy of thresholds obtained from psychometric functions. *Spatial Vision, 11*, 135–139.

Foster, D. H., Marín-Franch, I., Amano, K., & Nascimento, S. M. C. (2009). Approaching ideal observer efficiency in using color to retrieve information from natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision, 26*, B14–B24.

Foster, D. H., & Nascimento, S. M. C. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London Series B – Biological Sciences, 257*, 115–121.

Foster, D. H., Nascimento, S. M. C., Amano, K., Arend, L., Linnell, K. J., Nieves, J. L., et al. (2001). Parallel detection of violations of color constancy. *Proceedings of the National Academy of Sciences of the United States of America, 98*, 8151–8156.

Foster, D. H., & Snelgar, R. S. (1983). Test and field spectral sensitivities of colour mechanisms obtained on small white backgrounds: Action of unitary opponent-colour processes? *Vision Research, 23*, 787–797.

Funt, B., Ciurea, F., & McCann, J. (2004). Retinex in MATLAB™. *Journal of Electronic Imaging, 13*, 48–57.

Funt, B. V., & Drew, M. S. (1993). Color space analysis of mutual illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 15*, 1319–1326.

Gijsenij, A., Gevers, T., & van de Weijer, J. (2012). Improving color constancy by photometric edge weighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 34*, 918–929.

Giles, P. T. (2001). Remote sensing and cast shadows in mountainous terrain. *Photogrammetric Engineering and Remote Sensing, 67*, 833–839.

Goria, M. N., Leonenko, N. N., Mergel, V. V., & Novi Inverardi, P. L. (2005). A new class of random vector entropy estimators and its applications in testing statistical hypotheses. *Journal of Nonparametric Statistics, 17*, 277–297.

Gu, L., & Robles-Kelly, A. (2014). Shadow modelling based upon Rayleigh scattering and Mie theory. *Pattern Recognition Letters, 43*, 89–97.

Hagedorn, J., & D'Zmura, M. (2000). Color appearance of surfaces viewed through fog. *Perception, 29*, 1169–1184.

Heckman, G. M., Muday, J. A., & Schirillo, J. A. (2005). Chromatic shadow compatibility and cone-excitation ratios. *Journal of the Optical Society of America A – Optics Image Science and Vision, 22*, 401–415.

Hernández-Andrés, J., Romero, J., & Nieves, J. L. (2001). Color and spectral analysis of daylight in southern Europe. *Journal of the Optical Society of America A – Optics Image Science and Vision, 18*, 1325–1335.

Hubel, P. M. (2000). The perception of color at dawn and dusk. *Journal of Imaging Science and Technology, 44*, 371–375 (387).

Hurlbert, A., & Wolf, K. (2004). Color contrast: A contributory mechanism to color constancy. *Progress in Brain Research, 144*, 147–160.

Jiang, X., Schofield, A. J., & Wyatt, J. L. (2011). Shadow Detection based on Colour Segmentation and Estimated Illumination. In J. Hoey, S. McKenna, & E. Trucco (Eds.), *British Machine Vision Conference 2011* (pp. 87.1–87.11). BMVA Press.

Judd, D. B., MacAdam, D. L., & Wyszecki, G. (1964). Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America, 54*, 1031–1040.

Kamermans, M., Kraaij, D. A., & Spekreijse, H. (1998). The cone/horizontal cell network: A possible site for color constancy. *Visual Neuroscience, 15*, 787–797.

Kent, J. T., & Mardia, K. V. (2012). A geometric approach to projective shape and the cross ratio. *Biometrika, 99*, 833–849.

Kentridge, R. W., Heywood, C. A., & Cowey, A. (2004). Chromatic edges, surfaces and constancies in cerebral achromatopsia. *Neuropsychologia, 42*, 821–830.

Khang, B.-G., & Zaidi, Q. (2002). Cues and strategies for color constancy: Perceptual scission, image junctions and transformational color matching. *Vision Research, 42*, 211–226.

Kingdom, F. A. A., Beauce, C., & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception, 33*, 907–914.

Kozachenko, L. F., & Leonenko, N. N. (1987). Sample estimate of the entropy of a random vector. *Problems of Information Transmission, 23*, 95–101 (Tr. Problemy Peredachi Informatsii. Vol. 23, No. 2, pp. 9–16, 1987).

Kulikowski, J. J., Daugirdiene, A., Panorgias, A., Stanikunas, R., Vaitkevicius, H., & Murray, I. J. (2012). Systematic violations of von Kries rule reveal its limitations for explaining color and lightness constancy. *Journal of the Optical Society of America A – Optics Image Science and Vision, 29*, A275–A289.

Land, E. H. (1983). Recent advances in retinex theory and some implications for cortical computations: Color vision and the natural image. *Proceedings of the National Academy of Sciences of the United States of America, 80*, 5163–5169.

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America, 61*, 1–11.

Lee, R. J., Dawson, K. A., & Smithson, H. E. (2012). Slow updating of the achromatic point after a change in illumination. *Journal of Vision, 12*(1). http://dx.doi.org/10.1167/12.1.19. 19, 1–22.

Lee, T. Y., & Brainard, D. H. (2014). The effect of photometric and geometric context on photometric and geometric lightness effects. *Journal of Vision, 14*(1). http://dx.doi.org/10.1167/14.1.24. 24, 1–17.

Leek, E. C., Davitt, L. I., & Cristino, F. (2015). Implicit encoding of extrinsic object properties in stored representations mediating recognition: Evidence from shadow-specific repetition priming. *Vision Research, 108*, 49–55.

Linnell, K. J., & Foster, D. H. (1996). Dependence of relational colour constancy on the extraction of a transient signal. *Perception, 25*, 221–228.

Lovell, P. G., Tolhurst, D. J., Párraga, C. A., Baddeley, R., Leonards, U., Troscianko, J., & Troscianko, T. (2005). Stability of the color-opponent signals under changes of illuminant in natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision, 22*, 2060–2071.

Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences, 2*, 288–295.

Marín-Franch, I., & Foster, D. H. (2013). Estimating information from image colors: An application to digital cameras and natural scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 35*, 78–91.

Maybank, S. J. (1995). Probabilistic analysis of the application of the cross ratio to model-based vision: Misclassification. *International Journal of Computer Vision, 14*, 199–210.

Middleton, W. E. K. (1960). Bouguer, Lambert, and the theory of horizontal visibility. *Isis, 51*, 145–149.

Moons, T., Pauwels, E. J., Van Gool, L. J., & Oosterlinck, A. (1995). Foundations of semi-differential invariants. *International Journal of Computer Vision, 14*, 25–47.

Murray, I. J., Daugirdiene, A., Panorgias, A., Stanikunas, R., Kulikowski, J., & Kelly, J. M. F. (2014). Lightness constancy and its link with cone contrast. *Journal of the Optical Society of America A – Optics Image Science and Vision, 31*, A350–A356.

Nascimento, S. M. C., Amano, K., & Foster, D. H. (2016). Spatial distributions of local illumination color in natural scenes. *Vision Research, 120*, 39–44.

Nascimento, S. M. C., Ferreira, F. P., & Foster, D. H. (2002). Statistics of spatial cone-excitation ratios in natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision, 19*, 1484–1490.

Nascimento, S. M. C., & Foster, D. H. (1997). Detecting natural changes of cone-excitation ratios in simple and complex coloured images. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 264*, 1395–1402.

Nascimento, S. M. C., & Foster, D. H. (2000). Relational color constancy in achromatic and isoluminant images. *Journal of the Optical Society of America A – Optics Image Science and Vision, 17*, 225–231.

Nascimento, S. M. C., & Foster, D. H. (2001). Detecting changes of spatial cone-excitation ratios in dichoptic viewing. *Vision Research, 41*, 2601–2606.

Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., & Limperis, T. (1997). *Geometrical considerations and nomenclature for reflectance*. Washington, D.C.: Institute for Basic Standards, National Bureau of Standards.

Osorio, D., & Vorobyev, M. (2005). Photoreceptor spectral sensitivities in terrestrial animals: adaptations for luminance and colour vision. *Proceedings of the Royal Society B – Biological Sciences, 272*, 1745–1752.

Párraga, C. A., Brelstaff, G., Troscianko, T., & Moorehead, I. R. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision, 15*, 563–569.

Peyvandi, S., Nieves, J. L., & Gilchrist, A. (2013). On the information content along edges in trichromatic images. *21st Color and Imaging Conference 2013: Color Science and Engineering Systems, Technologies, and Applications*, CIC 2013 (pp. 236-239). Albuquerque, NM: IS&T and SID.

Pinto, P. D. A. (2004). *Chromatic properties of natural scenes and their influence on gaze position* (M. Phil. thesis). Manchester, UK: University of Manchester Institute of Science and Technology.

Plet, S., & Gerbino, W. (2001). Background articulation and relational colour constancy. *Color Research and Application, 26*, S201–S204.

Raath, N., & Woodward, S. (2008). *Unpublished MPhys project report*. School of Physics and Astronomy: University of Manchester, UK.

Reeves, A. J., Amano, K., & Foster, D. H. (2008). Color constancy: Phenomenal or projective? *Perception & Psychophysics, 70*, 219–228.

Rinner, O., & Gegenfurtner, K. R. (2002). Cone contributions to colour constancy. *Perception, 31*, 733–746.

Ripamonti, C., Bloj, M., Hauck, R., Mitha, K., Greenwald, S., Maloney, S. I., et al. (2004). Measurements of the effect of surface slant on perceived lightness. *Journal of Vision, 4*, 747–763.

Ripamonti, C., & Westland, S. (2003). Prediction of transparency perception based on cone-excitation ratios. *Journal of the Optical Society of America A – Optics Image Science and Vision, 20*, 1673–1680.

Romero, J., Luzón-González, R., Nieves, J. L., & Hernández-Andrés, J. (2011). Color changes in objects in natural scenes as a function of observation distance and weather conditions. *Applied Optics, 50*, F112–F120.

Ruderman, D. L., Cronin, T. W., & Chiao, C.-C. (1998). Statistics of cone responses to natural images: implications for visual coding. *Journal of the Optical Society of America A – Optics Image Science and Vision, 15*, 2036–2045.

Salvador, E., Cavallaro, A., & Ebrahimi, T. (2004). Cast shadow segmentation using invariant color features. *Computer Vision and Image Understanding, 95*, 238–259.

Schofield, A. J., Rock, P. B., & Georgeson, M. A. (2011). Sun and sky: Does human vision assume a mixture of point and diffuse illumination when interpreting shape-from-shading? *Vision Research, 51*, 2317–2330.

Sims, D. A., & Gamon, J. A. (2002). Relationships between leaf pigment content and spectral reflectance across a wide range of species, leaf structures and developmental stages. *Remote Sensing of Environment, 81*, 337–354.

Smithson, H., & Zaidi, Q. (2004). Colour constancy in context: Roles for local adaptation and levels of reference. *Journal of Vision, 4*, 693–710.

Stockman, A., & Sharpe, L. T. (2000). The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research, 40*, 1711–1737.

Stockman, A., Sharpe, L. T., & Fach, C. (1999). The spectral sensitivity of the human short-wavelength sensitive cones derived from thresholds and color matches. *Vision Research, 39*, 2901–2927.

Tappen, M. F., Freeman, W. T., & Adelson, E. H. (2005). Recovering intrinsic images from a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 27*, 1459–1472.

Tarr, M. J., Kersten, D., & Bülthoff, H. H. (1998). Why the visual recognition system might encode the effects of illumination. *Vision Research, 38*, 2259–2275.

Tokunaga, R., & Logvinenko, A. D. (2010). Material and lighting dimensions of object colour. *Vision Research, 50*, 1740–1747.

van Trigt, C. (2005). Illuminant-dependence of von Kries type quotients. *International Journal of Computer Vision, 61*, 5–30.

van Trigt, C. (2007). Linear models in color constancy theory. *Journal of the Optical Society of America A – Optics Image Science and Vision, 24*, 2684–2691.

von Kries, J. (1902). Theoretische Studien über die Umstimmung des Sehorgans. In: *Festschrift der Albrecht-Ludwigs-Universität* (pp. 145–158). Freiburg.

von Kries, J. (1905). Die Gesichtsempfindungen. In W. Nagel (Ed.), *Handbuch der Physiologie des Menschen* (pp. 211–212). Braunschweig: Vieweg und Sohn.

Vorobyev, M., & Osorio, D. (1998). Receptor noise as a determinant of colour thresholds. *Proceedings of the Royal Society of London Series B-Biological Sciences, 265*, 351–358.

Wachtler, T., Albright, T. D., & Sejnowski, T. J. (2001). Nonlocal interactions in color perception: Nonlinear processing of chromatic signals from remote inducers. *Vision Research, 41*, 1535–1546.

Wagemans, J., van Doorn, A. J., & Koenderink, J. J. (2010). *The shading cue in context. i-Perception, 1*, 159–178.

Warton, D. I., Wright, I. J., Falster, D. S., & Westoby, M. (2006). Bivariate line-fitting methods for allometry. *Biological Reviews, 81*, 259–291.

Werner, A. (2014). Spatial and temporal aspects of chromatic adaptation and their functional significance for colour constancy. *Vision Research, 104*, 80–89.

Westland, S., & Ripamonti, C. (2000). Invariant cone-excitation ratios may predict transparency. *Journal of the Optical Society of America A – Optics Image Science and Vision, 17*, 255–264.

Yang, Z., & Purves, D. (2003). Image/source statistics of surfaces in natural scenes. *Network-Computation in Neural Systems, 14*, 371–390.

Zaidi, Q. (1998). Identification of illuminant and object colors: Heuristic-based algorithms. *Journal of the Optical Society of America A – Optics Image Science and Vision, 15*, 1767–1776.

Zhou, W., Huang, G., Troy, A., & Cadenasso, M. L. (2009). Object-based land cover classification of shaded areas in high spatial resolution imagery of urban areas: A comparison study. *Remote Sensing of Environment, 113*, 1769–1777.

Żychaluk, K., & Foster, D. H. (2009). Model-free estimation of the psychometric function. *Attention, Perception, & Psychophysics, 71*, 1414–1425.