Identification of illuminant and object colors: heuristic-based algorithms

Qasim Zaidi

College of Optometry, State University of New York, 100 East 24th Street, New York, New York 10010

Received January 30, 1998; accepted February 24, 1998; revised manuscript received March 4, 1998

In everyday scenes, from perceived colors of objects and terrains, observers can simultaneously identify objects across illuminants and identify the nature of the light, e.g., as sunlight or cloudy. As a formal problem, identifying objects and illuminants from the color information provided by sensor responses is underdetermined. It is shown how the problem can be simplified considerably by the empirical result that chromaticities of sets of objects under one illuminant are approximately affine transformations of the chromaticities under spectrally different illuminants. Algorithms that use the affine nature of the correlation as a heuristic can identify objects of identical spectral reflectance across scenes lit simultaneously or successively by different illuminants. The relative chromaticities of the illuminants are estimated as part of the computation. Because information about objects and illuminants is useful in many different tasks, it would be more advantageous for the visual system to use such algorithms to extract both sorts of information from retinal signals than to discount either automatically at an early neural stage. [S0740-3232(98)01207-1]

OCIS codes: 150.0150, 100.0100, 330.0330, 330.1720, 330.4060.

1. INTRODUCTION

Identification of objects is an important function for the visual system to perform, and perceived colors of objects are useful in this task. Since the spectral composition of the light reflected from an object changes as a function of the spectral composition of the illuminant, it has frequently been suggested that the nervous system should discount the effect of the illuminant so that the colors of objects remain constant.[1–16] Many theories of human color constancy propose that early adaptation mechanisms[1–6] discard illuminant-related signals, while others suggest that these signals are discarded by higher-level adaptation[12] or simultaneous contrast.[7,13] The notion that the nervous system discards these signals has led to a large number of empirical studies that have attempted to measure the extent to which object colors are constant across illumination changes.[2,3,7,13–16]

The present paper, however, suggests that the problem for the visual system to solve is not to bring about stable color appearance under different illuminants by discounting them but to recognize that objects are indeed being viewed under different illuminants and to discover what the illuminant properties are.

The paper begins with observations that objects appear to be of systematically different colors under different illuminants. These transformations in the appearances of objects are information that the nervous system can potentially use in the identification of both objects and illuminants. Color photographs of natural scenes are used in this paper to demonstrate the systematic manner in which apparent colors of objects differ under different illuminants. Comparisons with gray-level versions of the same photographs reveal that observers can use differences in object colors to identify illuminants. Experimental studies have universally concluded that object colors as measured by asymmetric matching are not perfectly constant.[2,3,7,13–16] Instead of this failure of constancy being viewed as a limitation of the visual system, in this paper it is regarded as a design feature that allows the observer to extract information about the illuminant.[17]

Using the spectral reflectance functions of natural and manmade objects lit by different daylight illuminants, this paper shows that changes in the spectral composition of the illumination on a set of objects lead to affine transformations of the set of object chromaticities. Affine transformations have well-defined invariants,[18] and these invariants can be used to derive the transformation parameters, thus identifying identical objects across different sets of chromaticities. This is accomplished in two template-matching types of algorithms.

This paper also introduces the concept of the invariant "essential" color of an object. As a special case for the algorithms, if an observer has access to the essential color of even one identifiable object in a scene, the algorithms will recover the essential colors of other objects in the scene and of the illuminant.

2. ROLE OF COLOR IN PERCEIVING NATURAL SCENES

By looking at the same outdoor scene at various times of the day, through seasonal changes and under overcast or sunny skies, etc., one realizes that an object can appear to be of different colors under different natural illuminants. These shifts in color are useful cues from which an observer can infer the nature of the illuminant without having to see its source. Painters such as Monet and Corot used this fact to great effect in their paintings.

The chromaticity, brightness, color contrast, and brightness contrast of objects in a scene, coupled with memories of object colors in different lights, help the observer in making inferences about the nature of the illuminant. Compare color photographs of natural scenes[19,20] with gray-level versions of the same photographs (Fig. 1). Figure 1(a) shows shaded and sunny regions of an autumn scene. The observer easily identifies sunny regions as those areas where the foliage appears lighter and yellower. In this example, the perceived colors of objects clearly vary with illumination, conveying information about the illuminant spectrum. To attribute the yellowish appearance of the trees to the illuminant, however, the observer must assume a degree of similarity in the spectral reflectance functions of the foliage under the two illuminants. Because foliage that is perennially in the shade can be physically different from that exposed to sunlight, there is no guarantee that the reflectances are identical. But even if the observer's assumptions are imprecise, an approximate separation of perceived color into illuminant and object colors can be useful in many tasks, e.g., to a hiker seeking cooler or warmer trails. It could even be argued that in this case as in many others, object color invariance is of less utility than illuminant identification.

When we examine the gray-level version of Fig. 1(a), we might infer that knowledge of the geometry of shadows and junctions[21] would be sufficient to identify the lighter parts of the terrain as sunlit and that a memory of sunlight as yellow is not necessary to infer changes in the illuminant. The importance of color memory in inferring illuminant identify becomes clearer when the situation is compared with that in Fig. 1(b). Some parts of the boardwalk and the railing are lighter than other parts. Compared with the darker areas (gray-brown), the lighter parts have a blue-gray-green tinge. This tinge shows the light to be indirect and filtered through the green of the forest leaves. Although lit regions are easily identified in the gray-level versions of both Figs. 1(a) and 1(b), the identities of the illuminants can be inferred only if one

(a)

(b)

Fig. 1. (a) Sun dance, (b) community boardwalk, Port Projection. Photographs from Ketchum[19,20] scanned with an HP ScanJet 4P, converted to gray levels in Corel PhotoPal, and printed on a Tektronix Phaser IISD.

looks at the color photographs and sees the shift in the objects' colors.

The lower-level of contrast in Fig. 1(b) could be considered a clue to the lack of direct sunlight. However, the leaves in the lower-left part of Fig. 1(b) show that contrast alone can be misleading. In the gray-level version, this part of the scene has pale tones and high contrast, which could lead to the inference that it is receiving more light. In the color image, however, the light-yellow reflectance of the leaves is revealed to be the source of the lightness and higher contrast. The use of color information can thus supplement other strategies[21] for separating reflectance changes from illumination-caused changes.

Color photography is a nonlinear process, and the colors in Fig. 1 are only an approximation to those in the actual scenes. That the informal qualitative comparisons above are generally valid can be confirmed by examining real scenes. To summarize, a number of lessons can be learned by looking at natural scenes. The apparent colors of objects are different under different illuminants. Hence the apparent colors of objects contain information about both the surface reflectances and the illuminant spectra. Recognizing objects and illuminants can be useful in different tasks. Therefore an automatic early discounting of neural information concerning either would be suboptimal. The extraction of object and illuminant colors from a variegated scene seems to involve comparisons that can take place across retinal areas or time, when the set of reflectance functions under one illuminant can be assumed to be similar to a subset of reflectance functions under another illuminant. In addition, some inferences about object and illuminant colors seem to be referenced to memory colors. This paper will present algorithms that formalize such operations.

3. ESSENTIAL COLORS AS OBJECT AND ILLUMINANT INVARIANTS

For purposes of color vision, an object is characterized by its spectral reflectance function $[\theta_i(\lambda)]$, the proportion of light at each wavelength reflected by the object. An illuminant is characterized by its spectrum $[\Gamma_a(\lambda)]$, the energy of light at each wavelength. The quantum absorptions from an object $\theta_i(\lambda)$ under an illuminant $\Gamma_a(\lambda)$ are given by

$$\begin{aligned} S_{ia} &= \int s(\lambda)[\theta_i(\lambda) * \Gamma_a(\lambda)]\mathrm{d}\lambda, \\ M_{ia} &= \int m(\lambda)[\theta_i(\lambda) * \Gamma_a(\lambda)]\mathrm{d}\lambda, \\ L_{ia} &= \int l(\lambda)[\theta_i(\lambda) * \Gamma_a(\lambda)]\mathrm{d}\lambda, \end{aligned} \tag{1}$$

where $*$ is wavelength-by-wavelength multiplication and $s(\lambda)$, $m(\lambda)$, $l(\lambda)$ are the absorption spectra of the cones.

In general, object reflectances change over time, often as a result of exposure to light. For example, skin becomes more or less tan, flower colors change as they bloom, foliage turns a darker green from spring to summer, and clothes fade. However, to some extent, especially in the short run, surface reflectance functions can be assumed to be constant.[2,3] For any individual observer, to the extent that a surface reflectance function is invariant, so too are the scalar products of the surface reflectance function and the spectral sensitivity functions of that observer's cones:

$$\begin{aligned} S_i &= \int s(\lambda)\theta_i(\lambda)\mathrm{d}\lambda, \\ M_i &= \int m(\lambda)\theta_i(\lambda)\mathrm{d}\lambda, \\ L_i &= \int l(\lambda)\theta_i(\lambda)\mathrm{d}\lambda. \end{aligned} \tag{2}$$

Mathematically, these scalar products are identical to cone quantum catches from the object under a unit equal-energy illuminant $E(\lambda)$:

$$\begin{aligned} S_i &= \int s(\lambda)[\theta_i(\lambda) * E(\lambda)]\mathrm{d}\lambda, \\ M_i &= \int m(\lambda)[\theta_i(\lambda) * E(\lambda)]\mathrm{d}\lambda, \\ L_i &= \int l(\lambda)[\theta_i(\lambda) * E(\lambda)]\mathrm{d}\lambda. \end{aligned} \tag{3}$$

Hence for each observer a fixed function of the sensor responses from an object under an equal-energy illuminant can serve as the invariant "essential" color of that object.

For any individual observer, to the extent that an illuminant spectrum is invariant, so too are the cone quantum catches:

$$\begin{aligned} S_a &= \int s(\lambda)\Gamma_a(\lambda)\mathrm{d}\lambda, \\ M_a &= \int m(\lambda)\Gamma_a(\lambda)\mathrm{d}\lambda, \\ L_a &= \int l(\lambda)\Gamma_a(\lambda)\mathrm{d}\lambda. \end{aligned} \tag{4}$$

The cone quantum catches from an equal-energy illuminant are obviously proportional to the area under each cone spectral sensitivity curve.

Human cone spectral sensitivities are well approximated by the Smith–Pokorny[22] fundamentals. The outputs of the cones are combined into two opponent-color channels and a luminance channel. This serves to decorrelate cone signals, particularly the extremely high correlation between M and L signals.[23,24] The two color-opponent signals are represented by the axes of the MacLeod–Boynton[25] chromaticity diagram (Fig. 2). The horizontal axis represents $[\mathrm{L}/(\mathrm{L} + \mathrm{M})]$ and the vertical axis represents $[\mathrm{S}/(\mathrm{L} + \mathrm{M})]$. In this diagram an equal-energy light plots at (0.66, 0.016). For isolated lights the appearance of hues along the horizontal axis changes from greenish to reddish and along the vertical axis from yellowish to violet. For mnemonic reasons the coordinate axes will also be referred to as rg and yv.

The simulations in this paper used the spectra of the 170 natural and manmade objects that were measured by Vrhel *et al.*[26] The chromaticities of these objects under

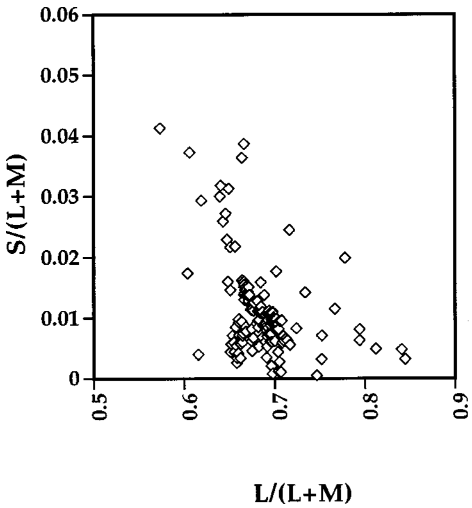

Fig. 2. MacLeod–Boynton chromaticities of 170 natural and manmade objects (Vrhel *et al.*[26]) in equal-energy light.

equal-energy light are shown in Fig. 2. These chromaticities can be considered to be the essential colors of these objects.

Similarly, for each observer the chromaticity of the illuminant relative to the chromaticity of the equal-energy illuminant can serve as the "essential" color of the illuminant.

The relation between color appearance and chromaticity is complex.[27] The formal algorithms in this paper operate exclusively on chromaticities, and the results of this paper are independent on any linking hypotheses between chromaticities and appearances. Where needed for informal purposes, the link will rely on the general assumptions that distinct chromaticities generally correlate with distinct appearances and that the direction of the difference in chromaticity space correlates with the perceived difference in dominant hues. For example, it is sufficient to assume that the different appearances of trees in shade and sunlight [Fig. 1(a)] correspond to differences in neural signals (chromaticities) and that comparisons between neural signals are potentially exploitable by object and illuminant identification algorithms.

4. RECOVERY OF OBJECT AND ILLUMINANT INVARIANTS

Sensor responses from an object can vary in a complex manner if the scene contains highlights, shadows, interreflections, etc. Therefore for formal analysis the situation has generally been reduced to that of flat matte surfaces under a spatially uniform diffuse illuminant.

The problem of color constancy has been defined as the recovery of invariant surface spectral reflectances $\theta_i(\lambda)$ from the sensor responses (S_{ia}, M_{ia}, L_{ia}).[28–35] This is an underdetermined problem that can be solved only under severely restrictive conditions and assumptions. This paper takes the approach that it is "neither correct nor helpful to see color constancy as a problem of measuring surface reflectance."[36] Instead the task is conceived of as providing surface descriptors that are the same across illuminants[36,37] and illuminant descriptors that are independent of the surfaces illuminated.

Essential colors would suffice as invariant descriptors, but it is clear from Eqs. (1)–(4) that recovering (S_i, M_i, L_i) and (S_a, M_a, L_a) is not possible from just the cone signals (S_{ia}, M_{ia}, L_{ia}). When certain additional assumptions are satisfied, it becomes possible to identify essential colors. However, in general, less complete recovery procedures may suffice. For many visual tasks the functions of color constancy can be accomplished by establishing correspondence among surfaces seen under different lights across either space or time. The observer can safely infer that corresponding surfaces have similar essential colors (or reflectances). Similarly, for illuminant identification purposes it is often sufficient to recover the color of a light relative to the color of another light, e.g., the identification of the yellower illuminant as sunlight in Fig. 1(a).

So that the problem will be tractable, the first algorithm presented in this paper will operate on the assumption that the visual system has access to sets of chromaticities from the same set of objects under two spectrally different illuminants. A second algorithm will treat the case in which one of the sets of objects is a subset of the other. The algorithms will attempt to match the same objects across illuminants, i.e., to show that particular pairs of chromaticities in the two sets belong to the same object. With the same operation, the algorithms will infer the difference in chromaticities between the illuminants, i.e., their relative colors. If one set of object chromaticities is essential, then the algorithms will derive the essential color of the other illuminant.

5. OBJECT CHROMATICITIES AND AFFINE TRANSFORMATIONS

The algorithms presented in this paper are essentially heuristic-based template matchers,[18,38,39] which start with sets of chromaticities under two different illuminants, calculated from the cone catches in Eq. (1). The task is to transform the chromaticities under one illuminant so as to align the two sets of chromaticities for the same objects on top of each and to recover illuminant information from the alignment procedure. The heuristic that simplifies this task is based on the systematic nature of the chromaticity shifts that occur when there is a change in the spectrum of the illumination.

In Fig. 3, in the left panel, each circle represents the L/(L + M) chromaticity coordinate of one object under equal-energy light (horizontal axis) versus the L/(L + M) chromaticity coordinate of the same object under zenith skylight[40] (vertical axis). The diagonal of unit slope is shown as a solid line. The energy spectrum of skylight is shown in the right panel. As Endler[41] has pointed out, light on an overcast day is close to an equal-energy light. The chromaticities of the illuminants are plotted at the center of the large dashed cross in the left and middle panels. The change between the illuminants leads to a highly correlated translation of object chromaticities along the L/(L + M) axis. Note that the line of shift in object chromaticities passes through the shift in illuminant chromaticities. The middle panel plots

S/(L + M) chromaticities of the same objects under the same two illuminants as in the left panel. The shift in object chromaticities is a correlated change of slope, i.e., a multiplicative shift. Another example of similar correlated chromaticity changes is shown in Fig. 4, which uses the same objects and compares equal-energy light to direct sunlight. Notice that under sunlight, object chromaticities along the S/(L + M) axis are shifted toward yellow relative to equal-energy light, whereas under skylight they are shifted toward the violet side. These shifts are consistent with the perceptual experience of shifts in object colors in natural scenes, for example in Fig. 1(a).

These correlated shifts in chromaticities are a consequence of correlated changes in all three classes of cone quantal absorptions[42–44] and the extremely high correlation between L- and M-cone quantum catches.[23,24] Correlated shifts in cone quantum catches have been demonstrated for a large variety of objects and illuminants[42–44] and are a consequence of integration within fairly broadband cone spectra. In fact, for all pairs of changes between phases of natural daylight,[40] simulations showed that, to a good approximation, illuminant spectrum changes lead to an additive change along the rg axis and a multiplicative one along yv.

To exploit the invariants of affine transformations, we can summarize these changes in the form

$$\begin{bmatrix} rg_{ia} \\ yv_{ia} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \sigma_{ab} \end{bmatrix} \begin{bmatrix} rg_{ib} \\ yv_{ib} \end{bmatrix} + \begin{bmatrix} \tau_{ab} \\ 0 \end{bmatrix}, \tag{5}$$

where i is the index for objects in the set, (rg_{ia}, yv_{ia}) and $(rg_{ib}, yv_{ib})$are the L/(L + M) and S/(L + M) chromaticities of the objects under illuminants A and B, and (τ_{ab}, σ_{ab}) describe the shift in the chromaticity of illuminant A relative to the chromaticity of illuminant B.

It is obvious from Eq. (5) that the chromaticities of a

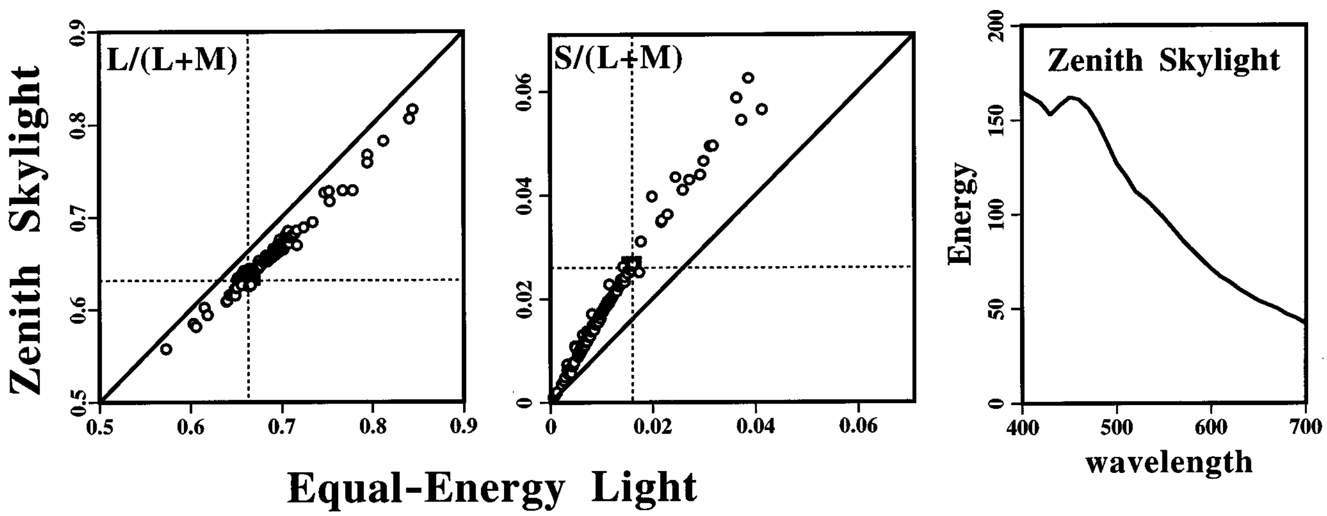

Fig. 3. Left, L/(L + M) chromaticities of 170 objects (Vrhel *et al.*[26]) in equal-energy light versus chromaticities of the same objects in zenith skylight; center, S/(L + M)chromaticities of the same objects under the same illuminants; right, energy spectrum of zenith skylight.

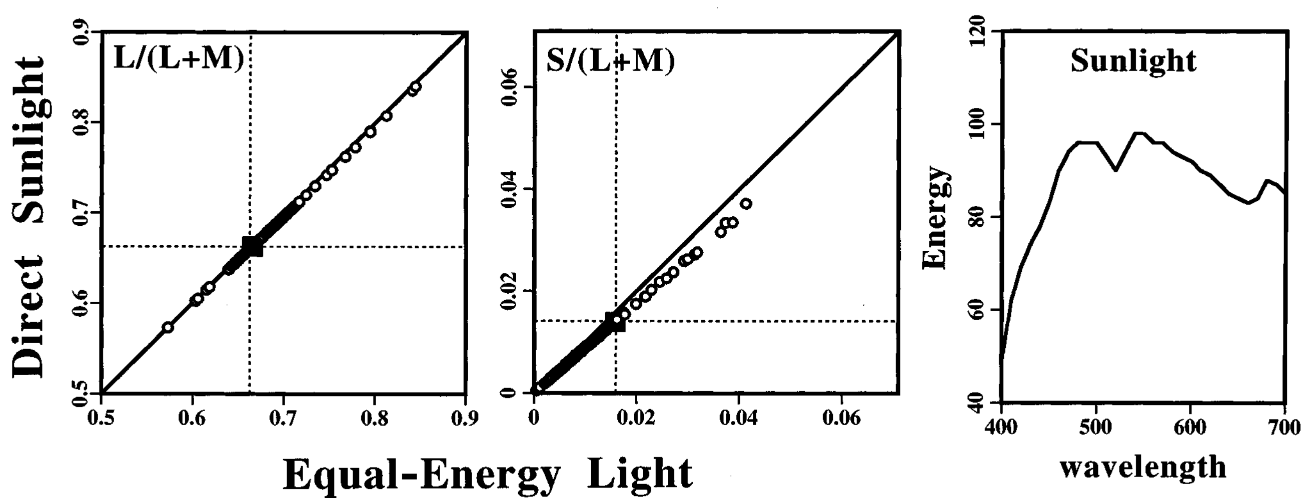

Fig. 4. Left, L/(L + M) chromaticities of 170 objects (Vrhel *et al.*[26]) in equal-energy light versus chromaticities of the same objects in direct sunlight; center, S/(L + M) chromaticities of the same objects under the same illuminants; right, energy spectrum of direct sunlight.

single identifiable object under two illuminants would be sufficient to derive the shift in all object chromaticities and the relation (τ, σ) between the chromaticities of the two illuminants. This result will be exploited in the two algorithms that follow.

6. IDENTIFICATION ALGORITHM FOR THE SAME SET OF OBJECTS UNDER TWO DIFFERENT ILLUMINANTS

In the case in which chromaticities from the same set of objects are available under two illuminants, the algorithm simply estimates τ_{ab}, the illuminant chromaticity shift along the rg axis, as the average difference between rg_{ia} and rg_{ib} for all i in the set. Similarly, σ_{ab}, the illuminant chromaticity scalar along the yv axis, is obtained from the average ratio of yv_{ia} to yv_{ib} for all i. These derivations are based on the property of affine transformations that the transform of the chromaticity that is at the center of gravity of the first set is always equal to the center of gravity of the transformed set. Calculation of the average rg difference is not a problem, so the obvious question is how to pick corresponding chromaticities from the two sets to calculate the average yv ratio. Excellent alignment between pairs of chromaticities could be obtained by exploiting more involved affine invariants as is done in computer object recognition.[18,38,39] However, a simpler strategy gives satisfactory results in this case. As shown in Figs. 3 and 4, the rank order of object chromaticities is generally preserved along each axis. Therefore in practice it is sufficient to rank order each set of chromaticities individually and to compute differences and ratios between pairs of the same rank order.

Figure 5 shows the result of applying the above algorithm to random samples of 50 chromaticities (from among the 170 objects) illuminated by equal-energy light (diamond) and each of the illuminants, whose chromaticities are shown as circles. These included five phases of natural sunlight,[40] a tungsten light, and a fluorescent light. Since the comparison illuminant was equal energy, essential illuminant chromaticities can be recovered by applying the shifts (τ_{aE}, σ_{aE}) to the chromaticity of the equal-energy spectrum. The crosses represent the chromaticities recovered by the algorithm for each of the illuminants. The crosses for each illuminant are close to the respective circles, indicating that despite the approximate heuristics used in the algorithm, illuminant shifts for a variety of illuminants can be recovered reliably. The algorithm works for any pair of illuminants, but in the absence of knowing the essential chromaticities for either illuminant, (τ_{ab}, σ_{ab}) will provide information only about the hue shift between the essential colors of the illuminants.

If two scenes contain the same objects, even in different spatial arrangements, through this algorithm the illuminant-caused chromaticity shift can be derived simply. If need be, one can then apply the shift to the chromaticities of individual objects under one illuminant to find objects with the same spectral reflectance functions under the second illuminant.

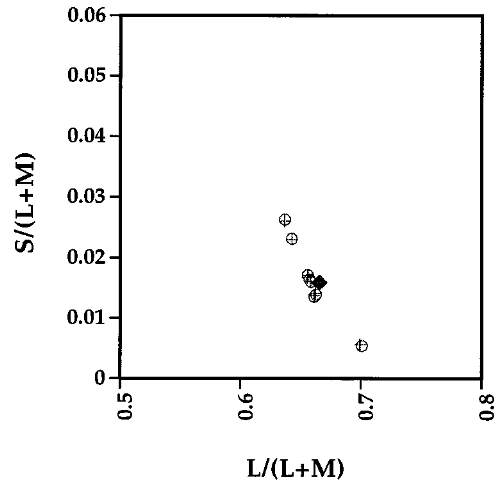

Fig. 5. Recovery of illuminant chromaticities by the algorithm. The diamond depicts equal-energy light, which was one of the illuminants on random samples of 50 objects. The circles represent the actual chromaticity of the other illuminant in each of the simulation experiments. The crosses represent the chromaticities recovered by the algorithm.

7. IDENTIFICATION ALGORITHM FOR A SUBSET OF OBJECTS TO BE ALIGNED TO A SUPERSET

The case in which only a subset of the objects seen under one illuminant is seen under the other seems not to have been treated before. The first algorithm cannot be used in this case because the corresponding pairs of chromaticities are not known. Solving the correspondence problem is more difficult in this case, because we need to estimate not only the shape change in the group of chromaticities but also the best position for aligning the smaller set on the bigger set.

A simple algorithm was written in MATLAB and used in the simulations described below. The algorithm rank orders the chromaticity pairs (rg_{ia}, yv_{ia}) for all i in the superset and (rg_{jb}, yv_{jb}) for all j in the subset, by the rg component. Then it picks (rg'_{jb}, yv'_{jb}) the smallest rg chromaticity from the subset, matches it to one chromaticity in the superset (rg'_{ia}, yv'_{ia}), and estimates the shifts in illuminant chromaticities by

$$(\tau_{ab}, \sigma_{ab}) = (rg'_{jb} - rg'_{ia}, yv'_{jb}/yv'_{ia}). \tag{6}$$

With Eq. (5), the pair (τ_{ab}, σ_{ab}) is then used to estimate the transformed chromaticities for all j in the subset. With squared distance along the two axes used as the metric, each transformed subset chromaticity is assumed to belong to the same object as the nearest superset chromaticity. The total squared error is then calculated between the transformed subset chromaticities and the nearest chromaticities in the superset. The above procedure is extremely fast, so the algorithm calculates the squared error for aligning (rg'_{jb}, yv'_{jb}) to each chromaticity in the superset. The alignment that minimizes the squared error is chosen as the best fit, and final estimates of τ_{ab} and σ_{ab} are calculated from the average differences and ratios, respectively, of all those pairs of subset and superset chromaticities that have been identified as belonging to the same object.

Figure 6 shows some results of using this algorithm. The crosses in the left panel represent the chromaticities of six of the Vrhel *et al.*[26] objects illuminated by equal-energy light. The circles in the middle panel show the chromaticities of a larger sample of 17 Vrhel objects, illuminated by zenith skylight. The objects in the left panel are all included in the set in the middle panel. The task of the algorithm is to find the best additive transform along L/(L + M) and the best multiplicative transform along S/(L + M) so that each cross can be matched to the circle that belongs to the same object. The crosses in the right panel are the result of applying the best transform found by the algorithm to the crosses in the left panel. For comparison, those circles from the center panel that belong to the same objects as the crosses are replotted in the right panel. The fit of the crosses to the circles is not perfect, but it is impressive that such a close fit can be achieved by an algorithm that can solve the correspondence problem for randomly arranged sets of chromaticities, uses only simple heuristics, and is extremely rapid in its operation. Figure 7 shows the results of applying the same algorithm to objects illuminated by direct sunlight and equal-energy light. These objects were the same as those represented in Fig. 6. The fit of crosses to circles in the right panel is even better than in Fig. 6, probably because the required shift in chromaticities is smaller.

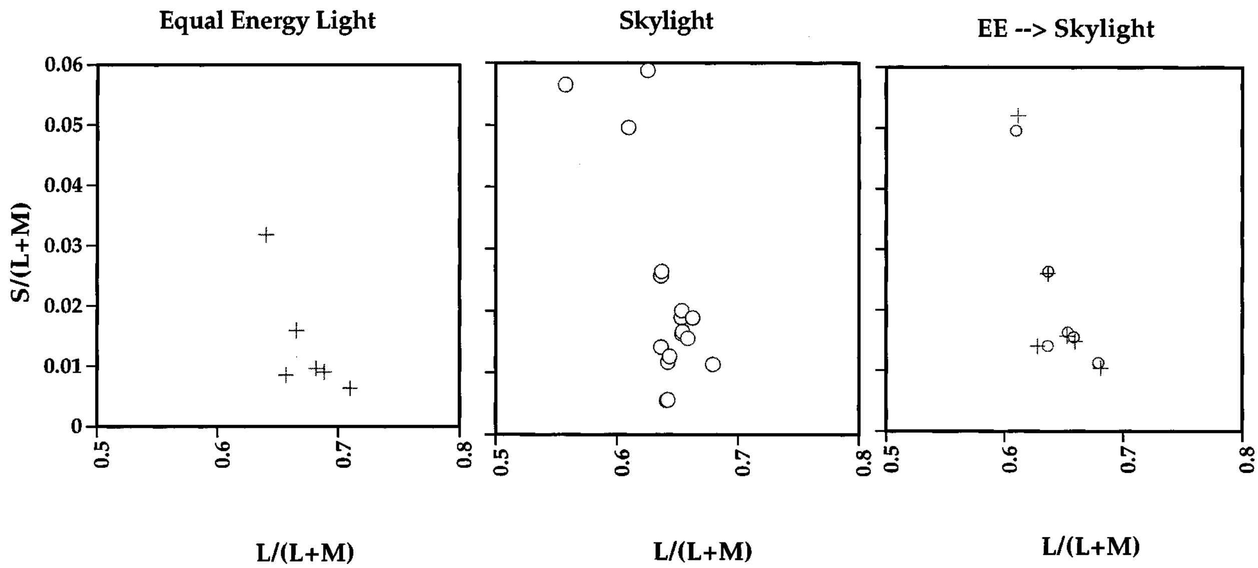

Fig. 6. Matching of objects by the algorithm despite illuminant-caused chromaticity shifts. Left, chromaticities of six objects from the Vrhel *et al.*[26] set lit by equal-energy light; center, chromaticities of 17 objects from the Vrhel set, lit by skylight. The six objects in the left panel are included in the 17. Right, crosses represent the results of applying to the crosses in the left panel the best affine transformation calculated by the algorithm. To show the accuracy of the matching procedure, circles for the same objects are replotted from the center panel.

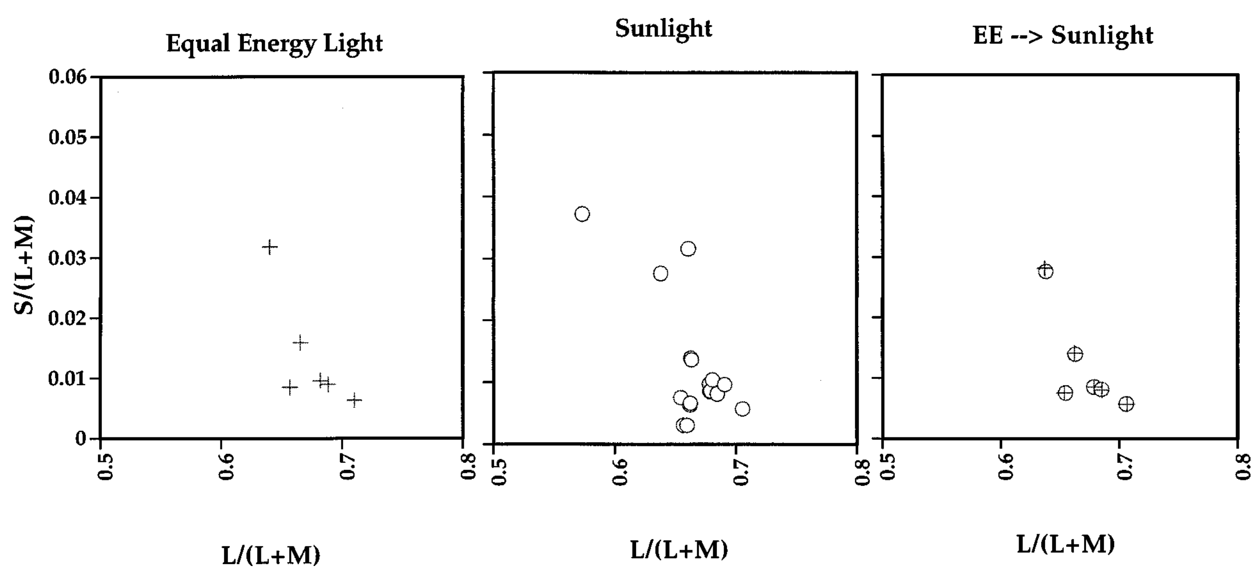

Fig. 7. Matching of objects by the algorithm despite illuminant-caused chromaticity shifts. Left, chromaticities of six objects from the Vrhel set lit by equal-energy light; center, chromaticities of 17 objects from the Vrhel set, lit by sunlight (the six objects in the left panel are included in the 17). Right, crosses represent the results of applying the best affine transformation calculated by the algorithm to the crosses in the left panel. To show the accuracy of the matching procedure, circles for the same objects are replotted from the center panel.

Any heuristic-based strategy can be "fooled" by carefully selected inputs, and the present algorithm relies on circumstances being close to satisfying affine transform assumptions. The algorithm has proven to be robust in a large number of simulations. For most subset–superset combinations, the squared error varied by up to a factor of 30 in choosing among alignments. Objects with the same spectral reflectance functions were almost always identified correctly across illuminants or identified with objects whose chromaticities were close to the veridical match. This robustness is due in part to the algorithm's search for matches between discrete object chromaticities. In additional simulations, when the algorithm was tested on small sets that did not share objects with the larger set, on the average, the squared error of the best alignment across illuminants was about five times larger than when subset objects were chosen from within the superset. The error was similarly large when the algorithm tried to align a nonintersecting smaller set to a larger set lit by the same illuminant.

The discussion above has glossed over some potential problems. Chromaticity space is not a metric space, so calculating distances between points (squared error) is not a strictly legitimate operation. It would be preferable to measure fits in a perceptual metric space in which errors in different color directions can be compared in terms of magnitude, but it would be a major task to find the metric space that corresponds to the particular adaptation state of the observer. Existing uniform color spaces are based on adaptation to spatially uniform fields, and the state of the visual system when adapted to variegated fields cannot be predicted from measurements made on spatially uniform fields.[45] It remains an empirical question whether the distortions introduced by using a squared norm in chromaticity space are tolerable. The stopping rule for the best fit is another issue that will be resolved as empirical experience accumulates. So far there has been no need to use more sophisticated procedures.

8. LEARNING TO IDENTIFY OBJECTS AND ILLUMINANTS

It would be extremely easy for a nervous system or a machine to learn to use these heuristic algorithms for matching objects across scenes and extracting relative colors of illuminants. All that is required is to learn from experience that illumination changes will shift object chromaticities along rg as a translation and along yv as a multiplicative scaling. This is much simpler than learning relative frequencies of occurrence of particular reflectance and illuminant functions, which is required for extracting object spectral reflectances by Bayesian estimation procedures.[34,35]

There is a claim in the literature that after viewing the same object under different illuminants, observers tend to remember its "reflectance."[46] Since the nervous system does not have access to object reflectances, it is probably more accurate to reword the above claim in terms of a memory for essential colors. If this type of memory exists for everyday objects and lights, it can be used in conjunction with the algorithms in this paper.

In the algorithms in this paper, if the correspondence of even one object is known under two unknown illuminants, the illuminant-caused shift (τ, σ) can be calculated exactly and applied to solve for the correspondence of other objects. A corollary of this property is that memorizing the essential colors of a few objects can help to estimate the essential color of the illuminant of a scene that contains even one of those objects. The estimated shift in turn can be applied to the other object chromaticities to estimate the essential colors of those objects. The same procedure can be followed if the essential color of the illuminant on a scene is known.

9. CONCLUSIONS

This paper began with demonstrations showing that perceived object colors change with illuminants and can thus be used to extract information about illuminants. The extent to which this is true about human observers needs to be investigated more systematically in a variety of controlled conditions. However, some published evidence does show that when observers are made to pay attention to the way things look, they perceive shifts in appearance caused by the illuminant.[47] In addition, observers are to some extent capable of identifying the surface properties of an object across two illuminants.[48] For empirical investigations the considerations presented in this paper imply that it is not profitable to measure the extent to which traditional concepts of color constancy hold if the purpose is to view any departure from perfect color constancy as a limitation of the visual system. Rather, it would be more profitable to examine the conditions under which observers can identify objects of similar surface reflectances under different illuminants despite the changes in appearance.

This paper presents two simple heuristic algorithms that match objects across illuminants and derive illuminant shifts purely from a pair of chromaticity sets. If such algorithms are implemented by the visual system, they can enable the observer to recognize when the same objects are being viewed under different illuminants despite there being a discernible shift in the colors of the objects. It remains to be investigated whether in the neural representation of color, affine invariants are preserved across spectral changes in the illuminant, which would enable the human visual system to use similar algorithms. One explanation for some recent "approximate-color-constancy" results[49] is that the visual system assumes that the center of gravity of chromaticity sets is preserved after an illumination change. The center of gravity is one of the invariants of affine transformations. Therefore the perceptual assumption about the center of gravity could be based on the stored heuristic that illumination changes on objects correspond to affine transformations on sets of chromaticities.

The most common models of color constancy in the literature invoke Von Kries–type receptor adaptation to equate neural signals from the same object across illuminants.[1] However, Von Kries scaling has been shown to be inadequate to describe the multistage nature of retinal and cortical adaptation mechanisms that can influence color appearance.[44,45] In addition, the assumptions that must be satisfied for successful automatic implementation of Von Kries scaling are unlikely to be true in most situations.[44] In variegated scenes, adaptation mechanisms do attenuate the differences between neural signals across illuminants, but in general the residual differences are greater than the limen of chromatic discrimination.[44] The algorithms in this paper can operate on the residual differences and are conceptually different from adaptation-based models of color constancy.

The approach in this paper is closest to machine-vision approaches based on canonical color gamuts, which also seek to provide surface color descriptors that are unaffected by changes in the illuminant.[36,37] The canonical color descriptor of an object is defined as the sensor responses to a patch under a fixed canonical illuminant.[36] The recovery procedure consists of constructing a prior canonical gamut for as many colors as possible under a single light and forming the convex hull of the gamut. For a set of surfaces under another illuminant, constraints on possible illuminants based on sensor responses are used to estimate a transform for the illuminant, which is applied to object sensor responses to obtain color descriptors, i.e., estimates of appearance under the canonical light. The adaptation of this procedure to two-dimensional chromaticity space has made it more robust to intensity variations and more efficient than in its three-dimensional form.[37] The present approach shares the advantages of operating in chromaticity space and has two additional advantages. First, the affine heuristic has been shown to be valid for a wide variety of illuminants and objects. Second, the template-matching nature of the algorithm obviates the need to store a canonical gamut and also does not require additional ad hoc heuristics such as maximizing the volume of the matched set of chromaticities. Note that Eqs. (2) and (3) show that if the canonical illuminant is specified to have an equal-energy spectrum, then the canonical color is an object invariant and is termed the essential color of the object.

Another mathematically sophisticated approach to color constancy has attempted to develop algorithms for estimating the spectra of reflectances and illuminants.[28–35] These algorithms are based on estimating each type of spectrum as a linear combination of a small number of fixed basis functions.[50–52] To function in general conditions, these algorithms require validated sets of basis functions, multiple views of many surfaces under different illuminants, and information about which surfaces under one illuminant correspond to which surfaces under the other illuminants. This approach has also been implemented in a Bayesian framework that incorporates prior marginal probability distributions of surface and illumination spectra.[34,35] The two algorithms presented in this paper not only are appreciably simpler and faster in their operation than the linear-basis algorithms, but they employ heuristics that are simpler to learn and validate, require only one pair of sets of chromaticities, and they solve the correspondence problem as part of extracting the color shift.

In addition, the second algorithm solves the correspondence problem and extracts illuminant shifts even when only a subset of objects is available under one of the illuminants. This is a direct consequence of the algorithmic search for discrete corresponding points that preserve affine invariants. As a consequence, this algorithm will be useful in many more situations than previous algorithms. For example, in Fig. 1(a) the trees in the sun are physically distinct and spatially separated from the trees in the shade. They have reflectances that are probably only a subset of the reflectances of the trees in the larger shaded regions, and corresponding reflectances are present in different spatial arrangements under the two illuminants.

It is possible that in visual processing, heuristics or priors that are based on correlations across relevant conditions are inherently more powerful than marginal priors. A similar conclusion can be drawn from the detection of motion, where detecting the orientation of spatiotemporal energy (i.e., correlations across time and space) is often more efficient than tracking the locations of identified features.[53]

ACKNOWLEDGMENTS

I thank Molly Aitken, Andrea Li, Fuzz Griffiths, and Mike Brill for extended discussions and help with this paper. A portion of this work was presented at the Association for Research in Vision and Ophthalmology Annual Meeting, Fort Lauderdale, Florida, 1997. This work was supported by National Eye Institute grant EY07556 to Q. Zaidi.

The author can be reached at the address on the title page or by e-mail: qz@cns.nyu.edu.

REFERENCES

1. H. E. Ives, "The relation between the color of the illuminant and the color of the illuminated object," Trans. Illum. Eng. Soc. **7**, 62–72 (1912). Reprinted in Color Res. Appl. **20**, 70–75 (1995).
2. E. Land and J. J. McCann, "Lightness and retinex theory," J. Opt. Soc. Am. **61**, 1–11 (1971).
3. E. Land, "Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image," Proc. Natl. Acad. Sci. USA **80**, 5163–5169 (1983).
4. G. West and M. H. Brill, "Necessary and sufficient conditions for Von Kries chromatic adaptation to give color constancy," J. Math. Biol. **15**, 249–258 (1982).
5. J. L. Dannemiller, "Computational approaches to color constancy: adaptive and ontogenetic considerations," Psychol. Rev. **96**, 255–266 (1989).
6. M. H. Brill, "Image segmentation by object color: a unifying framework and connection to color constancy," J. Opt. Soc. Am. A **7**, 2041–2049 (1990).
7. A. Valberg and B. Lange-Malecki, "'Colour constancy,' in Mondrian patterns: a partial cancellation of physical chromaticity shifts by simultaneous contrast," Vision Res. **30**, 371–380 (1990).
8. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: enhancing Von Kries adaptation via sensor transformations," in *Human Vision, Visual Processing, and Digital Display IV*, J. P. Allebach and B. E. Rogowitz, eds., Proc. SPIE **1913**, 473–484 (1993).
9. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: generalized diagonal transforms suffice," J. Opt. Soc. Am. A **11**, 3011–3019 (1994).
10. G. D. Finlayson and B. V. Funt, "Coefficient channels: derivation and relationship to other theoretical studies," Color Res. Appl. **21**, 87–96 (1996).

11. M. D. Fairchild and P. Lennie, "Chromatic adaptation to natural and incandescent illuminants," Vision Res. **32**, 2077–2085 (1992).
12. M. A. Webster and J. D. Mollon, "Colour constancy influenced by contrast adaptation," Nature **373**, 694–698 (1995).
13. J. Walraven, T. L. Benzshawel, B. E. Rogowitz, and M. P. Lucassen, "Testing the contrast explanation of color constancy," in *From Pigments to Perception*, A. Valberg and B. Lee, eds. (Plenum, New York, 1991), pp. 369–378.
14. H. Helson, D. Judd, and M. Warren, "Object-color changes from daylight to incandescent filament illumination," Illum. Eng. **47**, 221–233 (1952).
15. J. McCann, S. McKee, and T. Taylor, "Quantitative studies in retinex theory," Vision Res. **16**, 445–458 (1976).
16. D. H. Brainard, "Color constancy in the nearly natural image. 2. Achromatic loci," J. Opt. Soc. Am. A **15**, 307–325 (1998).
17. D. Jameson and L. Hurvich, "Essay concerning color constancy," Ann. Rev. Psychol. **40**, 1–22 (1989).
18. J. E. Hopcroft, D. P. Huttenlocher, and P. C. Wayner, "Affine invariants for model-based recognition," in *Geometric Invariance in Computer Vision*, J. L. Mundy and A. Zisserman, eds. (MIT Press, Cambridge, Mass., 1992), pp. 355–374.
19. R. G. Ketchum, *The Tongass: Alaska's Vanishing Rain Forest* (Aperture Foundation, New York, 1987).
20. R. G. Ketchum, *The Legacy of Wilderness: The Photographs of Robert Glenn Ketchum* (Aperture Foundation, New York, 1993).
21. P. Sinha and E. Adelson, "Recovering reflectance and illumination in a world of painted polyhedra," in *Proceedings of the Fourth International Conference on Computer Vision* (IEEE Computer Society Press, Los Alamitos, Calif., 1993), pp. 156–163.
22. V. C. Smith and J. Pokorny, "Spectral sensitivity of the foveal cone photopigments between 400 and 700 nm," Vision Res. **15**, 161–171 (1975).
23. G. Buchsbaum and A. Gottschalk, "Trichromacy, opponent colours coding, and optimum colour information transmission in the retina," Proc. R. Soc. London Ser. B **220**, 89–113 (1983).
24. Q. Zaidi, "Decorrelation of L- and M-cone signals," J. Opt. Soc. Am. A **14**, 3430–3431 (1997).
25. D. I. A. MacLeod and R. M. Boynton, "Chromaticity diagram showing cone excitation by stimuli of equal luminance," J. Opt. Soc. Am. **69**, 1183–1186 (1979).
26. M. Vrhel, R. Gershon, and L. S. Iwan, "Measurement and analysis of object reflectance spectra," Color Res. Appl. **19**, 4–9 (1994).
27. Q. Zaidi, "Color and brightness induction: from Mach bands to 3-D configurations," in *Color Vision: From Molecular Genetics to Perception*, K. Gegenfurtner and L. Sharpe, eds. (Cambridge U. Press, New York, to be published).
28. G. Buchsbaum, "A spatial processor model for object color perception," J. Franklin Inst. **310**, 1–26 (1980).
29. L. Maloney and B. Wandell, "Color constancy: a method for recovering surface spectral reflectance," J. Opt. Soc. Am. A **3**, 29–33 (1986).
30. M. D'Zmura and G. Iverson, "Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces," J. Opt. Soc. Am. A **10**, 2148–2165 (1993).
31. M. D'Zmura and G. Iverson, "Color constancy. II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces," J. Opt. Soc. Am. A **10**, 2166–2180 (1993).
32. M. D'Zmura and G. Iverson, "Color constancy. III. General linear recovery of spectral descriptions for lights and surfaces," J. Opt. Soc. Am. A **11**, 2389–2400 (1994).
33. S. Tominaga and B. A. Wandell, "Standard surface-reflectance model and illuminant estimation," J. Opt. Soc. Am. **6**, 576–584 (1989).
34. M. D'Zmura, G. Iverson, and B. Singer, "Probabilistic color constancy," in *Geometric Representations of Perceptual Phenomena. Papers in Honor of Tarow Indow's 70th Birthday*, R. D. Luce, M. D'Zmura, D. Hoffman, G. Iverson, and A. K. Romney, eds. (Erlbaum, Hillsdale, N.J., 1995), pp. 187–202.
35. D. H. Brainard and W. T. Freeman, "Bayesian color constancy," J. Opt. Soc. Am. A **14**, 1393–1411 (1997).
36. D. Forsyth, "A novel algorithm for color constancy," Int. J. Comput. Vision **30**, 5–36 (1990).
37. G. D. Finalyson, "Color in perspective," IEEE Trans. Pattern. Anal. Mach. Intell. **18**, 1034–1038 (1996).
38. D. P. Huttenlocher and S. Ullman, "Object recognition using alignment," in *Proceedings of the 1st International Conference on Computer Vision* (IEEE Computer Society Press, Los Alamitos, Calif., 1987), pp. 102–111.
39. D. P. Huttenlocher and S. Ullman, "Recognizing solid objects by alignment with an image," Int. J. Comput. Vision **5**, 195–212 (1990).
40. A. H. Taylor and G. P. Kerr, "The distribution of energy in the visible spectrum of daylight," J. Opt. Soc. Am. **31**, 3 (1941).
41. J. A. Endler, "The color of light in forests and its implications," Ecol. Monogr. **63**, 1–27 (1993).
42. J. L. Dannemiller, "Rank ordering of photoreceptor catches from objects are nearly illumination invariant," Vision Res. **33**, 131–137 (1993).
43. D. H. Foster and S. M. C. Nascimento, "Relational colour constancy from invariant cone-excitation ratios," Proc. R. Soc. London Ser. B **250**, 116–121 (1994).
44. Q. Zaidi, B. Spehar, and J. S. DeBonet, "Color constancy in variegated scenes: the role of low-level mechanisms in discounting illumination changes," J. Opt. Soc. Am. A **14**, 2608–2621 (1997).
45. Q. Zaidi, B. Spehar, and J. S. DeBonet, "Adaptation to textured chromatic fields," J. Opt. Soc. Am. A **15**, 23–32 (1998).
46. E. W. Jin and S. K. Shevell, "Color memory and color constancy," J. Opt. Soc. Am. A **13**, 1981–1991 (1996).
47. L. E. Arend, "How much does illuminant affect unattributed colors?" J. Opt. Soc. Am. **10**, 2134–2147 (1993).
48. L. E. Arend and A. Reeves, "Simultaneous color constancy," J. Opt. Soc. Am. A **3**, 1743–1751 (1986).
49. D. H. Brainard, W. A. Brunt, and J. S. Spiegle, "Color constancy in the nearly natural image. I. Asymmetric matches," J. Opt. Soc. Am. A **14**, 2091–2110 (1997).
50. J. Cohen, "Dependency of the spectral reflectance curves of the Munsell color chips," Psychon. Sci. **1**, 369–370 (1964).
51. D. B. Judd, D. L. MacAdam, and G. Wyszecki, "Spectral distribution of typical daylight as a function of correlated color temperature," J. Opt. Soc. Am. **54**, 1031–1040 (1964).
52. L. Maloney, "Evaluation of linear models of surface spectral reflectance with small numbers of parameters," J. Opt. Soc. Am. A **3**, 1673–1683 (1986).
53. J. S. DeBonet and Q. Zaidi, "Temporal and spatial frequency analysis of motion-energy and feature-tracking," Invest. Ophthalmol. Visual Sci. **S36**, 253 (1995).