Color Constancy

We would rather eat a banana that looks yellow rather than one that looks green, as the banana's color appearance carries reliable information about its ripeness. In general, color appearance is a useful percept because it provides reliable information about object identity and state. When we search for our car in a large parking lot, we rely on color to pick out likely candidates; our driver's licenses list the color of our eyes and hair for identification; we detect that a friend is embarrassed by the blush of his face.

For color appearance to provide useful information about objects, it must correlate with properties intrinsic to objects and be stable against transient features of the environment in which the objects are viewed, such as the illumination. This stability, which is provided in good measure by our visual systems, is called color constancy. That we have generally good constancy is consistent with everyday experience. We are content to refer to objects as having a well-defined color, and it is only rarely (e.g. when looking for our car in parking lot illuminated by sodium vapor lamps) that we observe large failures of constancy.

The Problem of Color Constancy

Vision obtains information about objects through the light reflected from them. If the reflected light were in one-to-one correspondence with physical object properties, then extracting stable object percepts would be straightforward. But the reflected light confounds properties of the illumination with those of the object. In the case of color, the relevant object property is its surface reflectance function $S(\lambda)$: the fraction of incident illuminant power that is returned to the eye at each wavelength λ. The relevant property of the illuminant is its spectral power distribution, $I(\lambda)$: the amount of power at each wavelength that arrives at the object. The spectrum reflected to the eye is thus $C(\lambda) = I(\lambda) S(\lambda)$. We call $C(\lambda)$ the color signal. Ambiguity arises because of the symmetric role played by $I(\lambda)$ and $S(\lambda)$ in the formation of the color signal. For example, a banana seen in skylight might reflect the same spectrum to the eye as grass under direct sunlight, because the effect of the illuminant change on the color signal can be counteracted by the change in surface reflectance. The perceptual challenge of color constancy is to make object color appearance stable against changes in $I(\lambda)$ while at the same time making it sensitive to changes in object reflectance $S(\lambda)$.

Empirical Observations

To what extent does the visual system actually stabilize object color in the face of illuminant changes? This has been studied with scaling and naming paradigms where observers describe the color appearance of objects seen under different illuminants, as well as with matching paradigms where observers adjust a test object seen under one illuminant to match the appearance of reference object seen under a different illuminant. A few generalizations may be drawn from a very large empirical literature. First, color appearance does vary somewhat when the illuminant is changed: color constancy is not perfect. Second, the variation in object appearance is small compared to what would be predicted for a visual system with no constancy.

That constancy is generally good is often characterized by a constancy index, which takes on a value of 0 for a visual system with no constancy and 1 for a visual system with perfect constancy. For natural viewing conditions when only the illuminant is changed, experimentally measured constancy indices are often in the range 0.8-0.9, and sometimes higher.

Constancy is not always good, however. For example, constancy fails for very simple scenes. Indeed, when a scene consists only of a single diffusely illuminated flat matte object, changes of really are perfectly confounded, and changes of illumination lead to large failures of constancy. More generally, the degree to which the visual system exhibits constancy depends critically on what is varied in the scene. Under natural viewing conditions, constancy tends to be very good if only the spectrum of the illuminant is varied. But if the surface reflectances of the other objects in the scene are covaried with the illuminant, constancy can be greatly reduced. For example, suppose the illuminant is shifted to have more short wavelength power and less long wavelength power. If at the same time, the surface reflectances of objects in the scene are shifted to compensate (i.e. to reflect less at short wavelengths and more at long wavelengths), color constancy is impaired. In laboratory studies of this manipulation, constancy indices drop into the range 0.2-0.4, a result that needs to be explained by any theory of constancy

Theories of Color Constancy

Theories of constancy should account for the general empirical observations described above. They should explain why constancy is often good, but also why it sometimes fails. Essentially all current theories share in common the general notion that visual processing of the color signal reflected from a single object is affected by the color signals reflected from the other objects in the scene. That is, our perception of object color is constructed by analyzing the reflected color signal relative to the rest of the retinal image.

Fundamentals of Color Vision

To understand theories of constancy, it is necessary to review a few fundamentals of human color vision. The color signal arriving at each retinal location is not represented completely. Rather, its spectrum is coded by the responses of three classes of light sensitive photoreceptors. These are referred to as the L, M, and S cones, where the letters are mnemonics for long-wavelength-sensitive, middle-wavelength-sensitive, and short-wavelength-sensitive. Each cone class is characterized by a spectral sensitivity that relates the cone's response to the intensity of incident light at each wavelength. The three classes cones differ in the region of the spectrum they are most sensitive to. Thus the information about color available to the brain consists of the responses r_L, r_M, and r_S of the L, M, and S cones at each image location.

Contrast Coding

The simplest theories of constancy postulate that the initial representation of the image is processed separately for each cone class, and that at each image location the cone responses are converted to a contrast representation. For the L cones, contrast is based on the difference between the overall L-cone response r_L and the L cone responses in its local neighborhood, and it expresses this difference relative to the magnitude of the neighboring responses. Let u_{L} represent the average of the L cone responses in a spatial neighborhood of an L cone whose response is r_L. Then the contrast is given by $c_L = (r_L - u_L)/u_{\mathrm{L}}$. Parallel expressions apply for the M and S cones.

Experiments that assess the visual system's response to spots flashed against spatially uniform backgrounds support the idea that cone responses are converted to a contrast representation early in the visual system. These experiments include measurements of appearance, of visual discrimination thresholds, and direct measurements of electrical activity in retinal ganglion cells.

How does contrast coding help explain color constancy? First consider a change in the overall intensity of the illuminant spectrum. This will increase the cone responses equally to the light reflected from every location in the scene.

Thus contrast, which depends on the ratio of cone responses across locations, is invariant to such illuminant changes.

When a change in the illuminant is not characterized as a simple scalar change in intensity, contrast is not guaranteed to be exactly invariant across the illuminant change. None-the-less, analyses of the light reflected from natural surfaces indicate that contrast computed separately for each cone class [i.e. c_L, c_m, c_S] is quite stable across naturally occurring illuminant changes, even those that are not pure intensity changes. When the objects in the scene are held fixed and only the spectrum of the illuminant is varied, theories based on local contrast coding do a good job of accounting for experimental color constancy data.

Challenges for Contrast Coding Models

Although contrast coding can explain constancy across illuminant changes, it also predicts strong failures of constancy when the reflectances of some of the objects in the scene are changed. Suppose we fix one test object of interest, and then vary the objects around it at the same time as the illuminant is changed. Manipulations of this sort can be arranged so that the local means [u_L, u_m , u_S] near the test object remain fixed across the illuminant change. In turn, this means that the contrasts of the test object [c_L, c_m, c_S] vary with the illuminant, and contrast theories thus predict large failures of constancy under these conditions. Although experimental tests do reveal that color constancy is reduced when this type of manipulation is performed, the reduction is not as great as contrast theories predict. In addition, contrast is a quantity computed independent of the three-dimensional of the scene being viewed. Experimental manipulations that affect perceived depth relations without changing cone contrasts have been shown to affect how color is perceived. Contrast-based theories are incomplete.

Theoretical Approaches and Considerations

The current theoretical challenge is to develop models that can account for the full range of experimentally measured performance, when stimulus manipulations include changes in the illuminant, in the reflectances of other objects in the scene, and in geometric aspects of the scene.

One approach to model development is mechanistic. The idea underlying mechanistic models is that the cone responses corresponding to each object are subjected to a series of transformations as they propagate from the retina to the cortex, and that the end result of these transformations is a neural representation that is more directly correlated with object reflectance than are the cone responses. The contrast coding idea reviewed above is an example of a mechanistic approach, and contrast coding is generally taken as the first stage of modern mechanistic theories. Current research along mechanistic lines seeks to incorporate additional stages. For example, most models postulate that signals from the separate cone classes are recombined at second stage opponent sites, which code differences between cone contrasts. Signals leaving the opponent sites are then further modified in an image-dependent fashion. For example, the overall magnitude of cone contrasts across the image are thought to control the gain applied to second-stage signals, and such processes have been shown to have the effect of further stabilizing color appearance across natural illuminant changes. To the extent that mechanistic models eventually account for color constancy across a wide range of experimental conditions, they have the attractive feature that they simultaneously provide an account of the chain of neural processing that underlies the color perception.

A second modeling approach is computational. Here the idea is to step back from the details of human vision and ask instead how to design an algorithm that had access to the cone responses could in principle achieve approximate color constancy. The algorithmic work lies within the domain of computer vision, and has led to methods that explicitly estimate the physical scene illuminant $I(\lambda)$ and surface reflectances $S(\lambda)$ from the spatial array of cone responses. Although the resulting algorithms do not have any necessary connection to human vision, they may be elaborated into models that predict how object colors will be perceived. For example, one can link algorithm and human performance by predicting that two objects viewed in different scenes and under different illuminants, but whose surface reflectances are estimated by the algorithm to be the same, will have also have the same color appearance. Modeling efforts of this sort have been successful at accounting for experimental measurements of object color appearance across a wide range of experimental conditions, and in particular at accounting for both successes and failures of constancy in a unified fashion.

Computational modeling is typically incorporates constraints derived from analysis of what surfaces and illuminants are occur in natural scenes. In the Bayesian approach to developing computational methods, such constraints are expressed explicitly as probability distributions. This in turn allows a precise statistical formulation of how to optimally estimate surface and illuminant properties from the cone responses. In the context of developing models for human vision, the motivation for incorporating optimality principles into the development of algorithms is that evolution has driven the visual system towards information processing optimized for natural scenes, so that a computational method that is optimized with respect to same set of scenes is likely to mimic human performance.

The mechanistic and computational approaches are complementary. Any computation that correctly predicts performance must have a mechanistic expression (since the human visual system implements it), and conversely any mechanistic formulation that describes performance must express a computation that does an excellent job of achieving color constancy for natural viewing (since we are only rarely surprised to see a fixed object appear a different color.) The current challenge is to elaborate both of these approaches to describe the full range of performance, and to bring them together into a single model that connects mechanism (including contrast coding), performance, and a computational understanding.

-- David H. Brainard

Cross References

Bayesian approach to perception; Color naming; Color perception; Color, Philosophical issues; Color perception, Physiological; Computational approach to perception; Constancy; Lightness constancy; Shape constancy; Size perception.

Further Readings

Brainard, D. H. (2004). Color Constancy. In L. Chalupa & J. Werner (Eds.), *The Visual Neurosciences* (Vol. 1, pp. 948-961). Cambridge, MA: MIT Press.

Mausfeld, R. and Heyer, D. (Eds.) (2003). *Colour Perception: Mind and the Physical World*. Oxford University Press: Oxford.

Shevell, S.K. and Kingdom, F.A.A. (2008). *Color in complex scenes.* Annual Review of Psychology. 59, 143-166.

Smithson, H.E. (2005). *Sensory, computational, and cognitive components of human* *color constancy*. Philosophical Transactions of the Royal Society of London B, 360*,* 1329-1346.