A68

Quasi-Mechanistic Accounts of Color Constancy

Donald I.A.MacLeod

Psychology, UCSD

LaJolla, CA 92093-0109

We can perceive the lightness and color of surfaces fairly accurately despite large variations in the intensity and spectral composition of the illumination. This color constancy is a form of visual adaptation in a general sense. But it is a different aspect of light adaptation--the associated change in visual sensitivity--that has been most completely documented in psychophysical and physiological experiments, and it is not yet clear to what extent the phenomena of constancy can be traced to the operation of known or plausible sensitivity regulating processes in the visual system. This paper will present an analysis of some of the evidence bearing on that issue.

Intuitively it seems reasonable to expect that changes in the intensity or color balance of the illuminant might be compensated by reciprocal adjustments of visual sensitivity, so that the resulting neural representation becomes independent of the illuminant; in order to accommodate chromatic variations of the illuminant, the sensitivity adjustments would have proceed more or less independently in different chromatic systems. Accordingly, experiments have shown a rough (though clearly not exact) independence of the different cone types in sensitivity regulation. Current quantitative treatments of color constancy, however, generally stress the inadequacy of this class of model: when the illumination on a scene changes in color, the cone excitations received from differently colored surfaces will change by different factors, so no sensitivity correction can be appropriate for all surfaces.

This argument would be a strong one if constancy were perfect. But constancy is not perfect, and I will suggest that models involving independent sensitivity corrections for the different cone signals, while clearly not completely adequate, do still provide a useful framework for understanding the constancy or inconstancy of color perception under changing conditions of observation: a better framework, perhaps, than the computationally perfect, but physiologically almost inconceivable, algorithms on which much attention is presently focussed.

Analysis of this problem is aided by considering a mathematically tractable idealization of the chromatic universe (including the observer), in which the logarithms of all wavelength-dependent quantities---the spectral reflectances of surfaces, the spectral energy distributions of illuminants, and the spectral sensitivities of the three cone types---are describable by second-order polynomials. Thus the cone receptor spectral sensitivities are approximated by Gaussians of suitable spectral peak and bandwidth; surfaces and illuminants can be of any color, with variation in spectral centroid, spectral dispersion or bandwidth and overall intensity or reflectance level. The cone excitations on which vision depends are each given by an expression involving nine parameters: three for the surface, three for the illuminant and three for the cone sensitivity function. This expression may be written as the product of three factors. First is a surface color factor, that depends on the proximity between the cone spectral peak and the spectral centroid (peak or minimum) of the surface reflectance function. Second, a surface-illuminant interaction factor depends on the proximity between the spectral centroids of the surface reflectance function and the illuminant energy distribution (but not on the cone parameters). The final factor depends on the proximity between the cone spectral peak and the illuminant spectral centroid; this factor is surface-independent and so can be exactly compensated by reciprocal adjustments of the sensitivity of each cone as illumination changes. The first factor (considered as a triple of numbers, one for each cone type) represents surface color independently of the spectral centroid of the illuminant. If, therefore, adaptive sensitivity adjustments do effectively hold the third factor constant, changing the spectral centroid of the illuminant can change only the second factor; this is independent of cone type, and hence affects only the effective intensity of the stimulus, and not its color as such.

To the extent that the simplifying assumptions are warranted, then, it turns out that a simple scaling of cone sensitivities can provide *perfect* color constancy, not in the strict sense that surfaces are visually unaffected by changing illumination, but in the sense that a constant appearance can be preserved by an adjustment of intensity (or overall reflectance level) alone. The need for such an intensity adjustment is evident: "When the light gets red, the reds get lighter." This holds not only at the physical level but also in visual experience, showing that the most notable failure of constancy in the sensitivity scaling model is not a failing of the model but a successfully represented characteristic of our visual processing. The *sufficiency* of the intensity adjustment is more surprising. It has been proposed that sensitivity regulation at the chromatic-opponent stage of processing would be useful in correcting the inadequacies of a cone-based constancy mechanism, but the present analysis suggests that this would actually be superfluous or harmful in providing compensation for changes in color balance in the illuminant.

Often, changes in illuminant color balance are not completely compensated by human observers in the way that a reciprocal sensitivity adjustment would predict. As others have noted, such partial constancy can be modelled by invoking sensitivity adjustments that fall some way short of reciprocity. Interactions between chromatic appearance and illumination or reflectance level have also been successfully accounted for with minor and independently supported refinements of the sensitivity regulation model: for instance, Helson's observation that light surfaces assume the illuminant color, and dark surfaces its complement, can be understood if the exponent of the response-intensity function for each cone increases with increasing cone excitation (Richards and Parks; Nayatani et al.), as brightness matching data (Whittle and Challands) indicate.

Changes in the illuminant bandwidth necessarily affect color rendering, as is clearest in the limiting case of monochromatic illumination. In our simplified chromatic universe the gamut generated by the surface-color factor shrinks with decreasing illuminant bandwidth. Again, this failure of color constancy appears to be characteristic of human observers, as well as of the model. Recent evidence, however, (Brown and MacLeod, these proceedings) indicates that such a physically constricted gamut can be perceptually enlarged, providing some degree of compensation for changes in illuminant bandwidth. This effect helps to explain the surprisingly good color rendering properties of predominantly monochromatic illumination (Helson; Boynton et al.) It can not be due to sensitivity adjustments at the cone level. The simplest physiological basis for it would be sensitivity regulation in cells with an opponent organization, but there are indications that this is too simple: interactions between color-opponent and non-opponent signals are needed to account for color appearance in these situations.

However successful they may be, models of constancy that rely entirely on low-level sensitivity adjustments are at best seriously incomplete. Several types of experiment have shown an important role for central processes in determining color appearance. Moreover, it is not clear physiologically how sensitivities could be appropriately set, since known regulating processes are too rapid and too local to average over a whole scene. Sensitivity adjustments that are rapid, local and reciprocal threaten to obliterate all perceived contrast; but near edges, involuntary eye movements will generate contrast-dependent signals, and these may incorporate the corrections on which constancy depends. This account, however, leaves unanswered the essential question of how such edge contrast signals are used to reconstruct perceived lightness and color across the whole field.

Supported by NIH Grant EY 01711