Reflections on colour constancy

Karl R. Gegenfurtner

A study of colour perception shows that, when assigning colour to objects, the seeing brain takes into account subtle reflections of light between the surfaces in a scene.

For more than two centuries, scientists and artists have come up with a range of ways to demonstrate that the wavelength composition over a whole scene can affect how we perceive the colour of the individual parts of that scene. The proportion of light of each wavelength reflected from an object can be highly beneficial in detecting or recognizing objects[1]. However, to make use of that invariant, the visual system somehow has to discount the illuminating light, which can vary quite drastically.

This process is usually called 'colour constancy', and the degree to which it is shown depends on many factors. Some of them occur at the early stages of sensory processing[2], such as local colour contrast, whereas others (for instance, colour memory) occur at higher cognitive levels[3]. Most computational schemes for achieving colour constancy try to decompose the overall light reaching the eye into one component that is due to the illuminant, and a second component due to the reflectance.

However, the physics of light is more complicated than the simple reflection of light from a surface into the eye, as if looking at a photograph. In a three-dimensional world, some light is reflected from one surface, but it then bounces to yet another surface from which it is reflected into the eye (Fig. 1). And so on. These indirect reflections are called 'inter-reflections', and are of especial interest to those involved in computer graphics and computer vision[4]. For example, computer simulations of indoor scenes appear more realistic when inter-reflections are taken into account. Indeed, many recent advances in computer graphics are due to the discovery of efficient algorithms to calculate the effects of all such multiple bounces of light among the vast number of surfaces typically contained in a scene.

Reporting on page 877 of this issue, Bloj *et al.*[5] tested whether inter-reflections are taken into account by the human visual system when perceiving the colour of surfaces. To do this, they presented subjects with a folded card, with its two sides opening towards the observer (Fig. 2). One side of the card was painted magenta, the other white. The magenta half reflects light only in part of the spectrum, some of which reaches the eye directly, and some of which is reflected onto the white side of the card. As a result, the light reflected from the white side is no longer balanced over all wavelengths, and it should appear a quite saturated magenta under the conditions of the experiment. Surprisingly, subjects reported that it appeared only slightly pinkish, suggesting that the visual system discounts the effect of inter-reflections to compute surface reflectance. In other words, the colour we perceive is influenced not only by the two-dimensional image of an object projected onto the retina, but also by our perception of the object's three-dimensional shape.

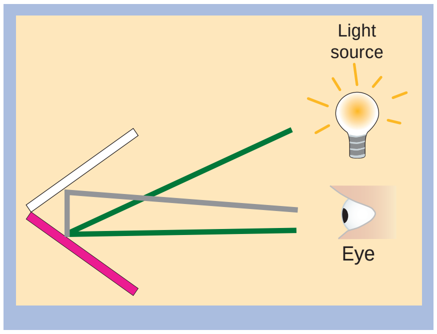

Figure 1 **Illustration of inter-reflections. Part of the light is being scattered directly into the eye of the observer (green line). Other light rays enter the eye after being reflected from the magenta side of the card onto the white side (grey line). The magenta is 'bleeding' into the white.**

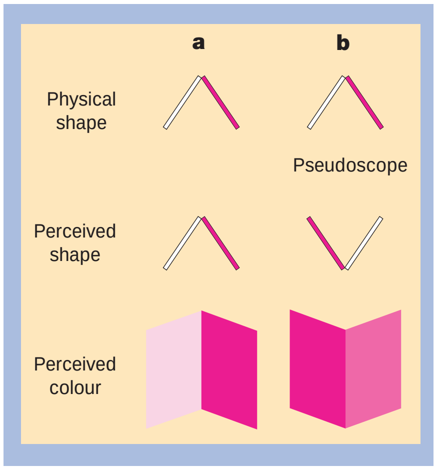

Figure 2 **The experiment done by Bloj *et al.*[5]. Observers viewed a folded card opening towards them.** a, **Under normal viewing, observers discount the effect of inter-reflections.** b, **When seen through a pseudoscope, the inter-reflections are interpreted as a change in surface colour because they are incompatible with the percept of the bent-away card. The magnitude of the effect has been emphasized for clarity.**

The strength of Bloj and colleagues' paper lies in the way they controlled for potential confounding factors. They used a special device called a pseudoscope, which optically inverts the depths of all objects in the scene. This means that the card will appear to open away from the observer when seen through the pseudoscope (Fig. 2b). This leads to an identical two-dimensional distribution of image intensities over the visual field. Perceptually, however, there is no longer the possibility for inter-reflections between the two surfaces, because they no longer seem to be pointing towards one another.

In this way, the physical stimulus was unchanged, so the inter-reflections and their effect were still present. But they could no longer be *interpreted* as such by the brain. Instead, the bias in the wavelength distribution of the card seemed to be due to its surface reflectance, and, as a result, it was perceived to have a deep pinkish hue. So the results show not only that three-dimensional shape can affect colour perception, but also that the seeing brain indeed knows the physics of light inter-reflections.

Similar phenomena are well known in our perception of brightness. When dealing with black and white scenes, the effects of shape, reflectance and illumination are basically interchangeable, presenting the visual system with an infinite number of possibilities for interpreting such scenes[6]. Yet we perceive the world as being quite stable. The solution seems to be that the most likely percept is calculated, so changes in the perception of three-dimensional shape can easily (and profoundly) affect our perception of brightness[7].

The question is, then, how much knowledge of the physics of light has the visual system acquired? It seems to know even the very subtle aspects, which may have only a very limited function when viewing natural outdoor scenes. For example, the geometric information contained in highlights — generated by mirror-like reflections from glossy surfaces — can affect the perception of surface curvature[8]. In their experimental set-up, Bloj *et al.*[5] chose conditions that emphasized the role of inter-reflections. The white side of the card almost lay in the shadow, so it received very little direct light; most light came from the magenta card. In typical real-life situations, however, a much smaller part of the light reaching the eye is due to inter-reflections.

A tremendous amount of psychophysical characterization has been done on visual 'illusions', such as those mentioned above, yet we know hardly anything about the brain processes that underlie such phenomena. The analysis of colour and other elementary features such as motion or orientation is often assumed to occur at an early stage of sensory processing in the visual system, whereas information about surfaces and objects is thought to be extracted during the later stages. Bloj and colleagues' results are difficult to reconcile with this simplified idea. One has to differentiate between the different uses of colour information in visual perception. Colour might be used at an early stage to segment objects from one another and from the background, but at that stage there is no need to assign a definitive colour to each object. Higher visual areas (such as the infero-temporal cortex, which is known to be important for object recognition[9]) also show a large degree of colour selectivity[10], and Bloj and colleagues' results support the idea that the colours we perceive may be determined at a rather late stage of visual processing. ■

Karl R. Gegenfurtner is at the Max-Planck-Institut für Biologische Kybernetik, Spemannstrasse 38, 72076 Tübingen, Germany.
e-mail: karl@kyb.tuebingen.mpg.de

1. Mollon, J. D. & Jordan, G. *J. Exp. Biol.* **146,** 21–38 (1988).
2. Kraft, J. M. & Brainard, D. H. *Proc. Natl Acad. Sci. USA* **96,** 307–312 (1999).
3. Troost, J. M. & de Weert, C. M. *Percept. Psychophys.* **50,** 591–602 (1991).
4. Foley, J. D., van Dam, A., Feiner, S. K. & Hughes, J. F. *Computer Graphics: Principles and Practice* (Addison–Wesley, Reading, Massachusetts, 1990).
5. Bloj, M. G., Kersten, D. & Hurlbert, A. C. *Nature* **402,** 877–879 (1999).
6. Adelson, E. H. & Pentland, A. P. in *Perception as Bayesian Inference* (eds Knill, D. C. & Richards, W.) 409–423 (Cambridge Univ. Press, New York, 1996).
7. Knill, D. C. & Kersten, D. *Nature* **351,** 228–230 (1991).
8. Blake, A. & Bülthoff, H. H. *Nature* **343,** 165–168 (1990).
9. Logothetis, N. K. & Sheinberg, D. L. *Annu. Rev. Neurosci.* **19,** 577–621 (1996).
10. Komatsu, H., Ideura, Y., Kaji, S. & Yamane, S. *J. Neurosci.* **12,** 408–424 (1992).

Photography

Making every photon count

Richard Hailstone

More than 150 years after the invention of photography, the underlying solid-state photochemistry is still being investigated and improved. One aspect of continual importance is the efficiency of the image-recording stage, which strongly affects the sensitivity of the photographic emulsion. Photons falling on a silver halide microcrystal in the emulsion cause a halide ion (such as chloride or bromide) to lose an electron, which is then captured by a silver ion, converting it into a neutral silver atom. The build-up of silver atoms gives rise to silver clusters: the 'latent image', which can be developed after several more steps into a print[1]. In reality, the halide ions generate pairs of electrons and holes (the positively charged counterparts of electrons), and it is generally accepted that recombination of photogenerated electrons and holes is a major loss process[2,3]. On page 865 of this issue, Belloni *et al.*[4] present an ingenious way of overcoming this recombination process by doping the silver halide with formate ions, which leads to an increase in the number of silver atoms generated per photon absorbed.

Image-forming reactions in silver halide microcrystals compete with various electron-loss processes, the dominant one being recombination. Overall, many of the photogenerated electrons are lost and the sensitivity of the silver halide emulsion is low. The sensitivity can be greatly improved by chemical sensitization, which is how all modern films and print materials are prepared. During chemical sensitization, reagents containing labile sulphur and gold atoms are used at low concentrations (parts per million or lower) to produce silver–gold sulphide clusters on the surface of the microcrystal. These clusters appear to deepen the existing traps for electrons, thereby reducing recombination[5–7].

During film development the silver halide microcrystal is converted to metallic silver through an electrochemical reduction reaction catalysed by the latent image. A microcrystal can only be developed when the latent image contains three or more photogenerated silver atoms[8]. In the theoretical limit, one photon should produce one silver atom. So the most sensitive film imaginable will only require three photons per microcrystal to create a latent image. Despite the advantages of sulphur-plus-gold sensitization there is still some recombination, and three to ten times as many photons are needed[9,10]. So reaching this theoretical limit requires an additional mechanism to eliminate the remaining recombination. In the past this has been done by reduction sensitization, which chemically creates silver clusters on the surface of the microcrystal. These silver clusters have different properties from the photogenerated silver clusters in that they trap and destroy holes through an oxidation process at the silver cluster[11]. In model systems (but not in commercial films) reduction sensitization, in addition to the sulphur-plus-gold sensitization, allows one to reach the theoretical limit[9,10].

Unfortunately a significant level of fog is often a byproduct of the combined sensitization processes. Fog is the result of latent-image formation in the absence of light exposure and, when present, reduces the dynamic range of the material for image capture. In addition, the stability of chemical reduction is often low, so it is not usually a viable option for photographic manufacturers. Placing the chemically produced silver clusters inside the microcrystal could eliminate the fog problem and presumably would stabilize the sensitization. But this may change the electronic properties of the silver clusters, because some experiments indicate that they have electron-trapping properties that would actually decrease the sensitivity[12].

Efficiency improvements are important to photographic manufacturers because efficiency is a factor that determines the film speed. Light absorption also determines film speed and this aspect is controlled by the size of the microcrystal. This is the usual way in which different film speeds are obtained.

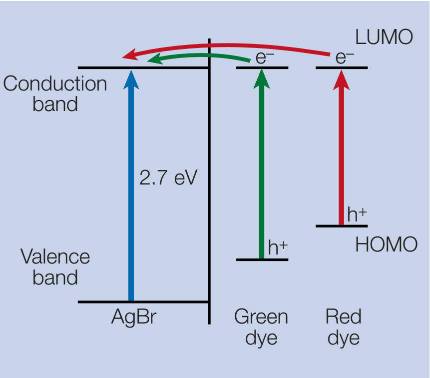

Figure 1 Colour photography requires three layers of silver halide emulsion sensitive to blue, red and green wavelengths of light. The effect of adding a green- or red-absorbing dye to AgBr (which absorbs blue light) is shown here. Absorption involves movement of an electron from the highest occupied molecular orbital (HOMO) to the lowest unoccupied molecular orbital (LUMO), leaving a hole in the HOMO level. The photoinduced electron transfer occurs between the LUMO and the AgBr conduction band, leaving behind an oxidized-dye radical. The LUMO levels for both the green-absorbing dye and the red-absorbing dye are positioned at the conduction-band minimum so that efficient electron transfer can occur. Thermal energy can be used to move the hole into the valence band, thereby preventing recombination with an electron. Belloni *et al.*[4] show that formate doping increases the sensitivity of green-dyed AgBr. Improving the overall efficiency of colour films might be possible, if similar increases in sensitivity can be shown for red-dyed emulsions.