Fred Kingdom's comments on Alan Gilchrist's chapter

file:///Volumes/LOP/brlt/index.htm

Just what belongs to what? The anchoring model applied to grating induction

One of the most interesting parts of Alan's talk was his exposition of how the anchoring model explains the 'standard' simultaneous lightness contrast (SLC) display illustrated below in Figure 1a. I would like to explore how the anchoring model would apply to another form of SLC, the grating induction (GI) stimulus illustrated in Figure 1b. The GI stimulus differs in important ways from the standard stimulus, but because its illusory lightness variations are believed to me mediated by the same mechanisms that are responsible for those in the standard display (e.g. see McCourt's commentary on Gilchrist's reply to these comments), it provides a good test of the generality of the anchoring model. How then does the anchoring model explain SLC in the standard display ? The idea is that the stimulus is divided into two local, and one global perceptual framework, with the lightness of each test patch being a weighted average of two lightness values, one computed from the local, and one from the global framework. The global framework comprises the stimulus as a whole, whereas the local frameworks are the two surrounds with their respective test patches, as illustrated in Figure 1a. Within each framework the anchor is the highest luminance, which is assigned white, and all other regions within the framework are assigned shades of grey that depend upon their luminance ratio with respect to the anchor. Within the global framework, the lighter of the two surrounds is the highest luminance and therefore assigned white, and the two test patches are assigned the same lightness values since their luminance ratios to the lighter surround are the same. However, for the two local frameworks the lightness assignments are different. In the case of the patch on the dark surround the patch is the anchor and is assigned white, whereas in the case of the grey patch on the light surround the surround is the anchor, and the patch is assigned a mid-grey relative to it. The resulting patch lightnesses are the averages of the global values (both patches equal and mid-grey), and the local values (patch on dark surround white, patch on light surround mid-grey). The resulting average value is higher for the patch on the dark surround than for the patch on the light surround, in accordance with the percept. Thus although the global framework influences the lightnesses of the patches, it is the operation of the two local frameworks that produces the difference in lightness between the patches, and hence the illusion we refer to as SLC.

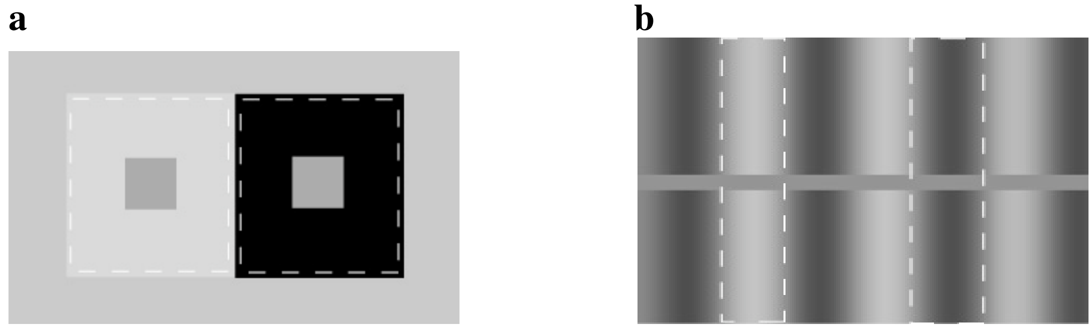

Figure 1. *(a) Standard simultaneous contrast display. The dashed lines show the two local perceptual frameworks of the anchoring model. According to the model, anchoring takes place within each framework, and it is this that gives rise to the illusory differences in the lightness of the two patches (see text for details). (b) Anchoring model as applied to grating induction. The local perceptual frameworks that result in the illusion are presumably the bright and dark bars of the inducing grating, together with those parts of the test stripe that pass through each. Two of these are delineated by the dashed lines.*

Turning now to grating induction (GI), the difference between GI and the standard SLC stimulus is twofold: the surround varies continuously in luminance and a single test region runs continuously through the surround. How would the anchoring model apply to GI? The two local frameworks must presumably be the dark and light inducing half-cycles, together with those parts of the test stripe that pass through each, as illustrated in Figure 1b. What strikes me about this putative framework configuration is just how arbitrary it is. One could just as plausibly argue that the inducing grating constitutes one perceptual framework, and the test stripe another, in which case the anchoring model would predict no lightness induction at all. To reinforce the point, consider the version of GI shown in Figure 2. When viewed stereoscopically the inducing pattern, now a cylinder, appears in a different depth plane from the test stripe. Cylinder and stripe are now unmistakably distinct objects, yet the induction remains. It is surely a stretch of the imagination to suppose that the two critical perceptual groupings in this figure are first the left half of the test stripe plus dark-shaded part of the cylinder, and second the right half of the test stripe plus light-shaded part of the cylinder, given the obvious alternative of cylinder and stripe. In short the anchoring model fails because the perceptual groupings required to make it work are simply implausible. Of course one can always argue that the perceptual groupings must be those that predict the illusion according to the model, but then the whole argument becomes circular.

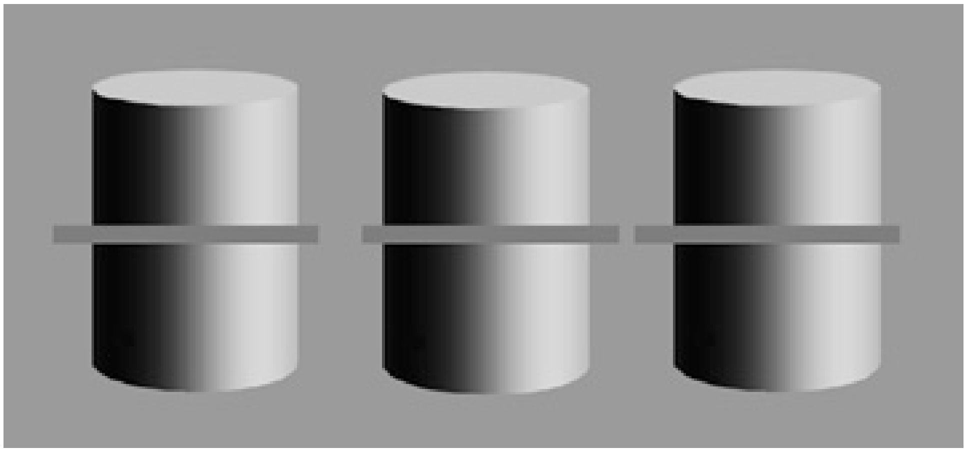

Figure 2. *Cylinder version of grating induction. When free-fused, one of the horizontal test stripes appears in front of the cylinder; lightness induction is still marked, even though cylinder and stripe are clearly different objects.*

I am not suggesting that anchoring, if indeed it operates in part locally, plays no role in producing illusions such as SLC. Rather, I assert that it is unlikely to be the main cause, because SLC occurs under conditions where the anchoring model cannot plausibly apply. Moreover, the evidence that SLC and GI are primarily a result of the operation of contrast-sensitive mechanisms provides a much more parsimonious explanation (see chapters by Kingdom and Blakeslee & McCourt). Having said that, there is indeed something attractive about the idea that anchoring is in part a local process, especially for those of us who believe in the importance of local processes in brightness/lightness perception. What is needed to convince an anchoring-model skeptic like me is some *independent* piece of behavioural evidence that anchoring operates locally. In other words, rather than showing that a model which incorporates anchoring on a local basis can be made to account for a range of lightness illusions, first identify a unique behavioural signature for anchoring, for example one of its temporal properties, and then show that it operates locally, and not just globally.

Gilchrist and intrinsic image models

Alan showed a number of compelling demonstrations in his talk. Some were in support of the anchoring model, while others served as ammunition against other models of lightness perception. Here I consider one of the demonstrations in the latter category, the one Alan used to argue against 'intrinsic image' models of lightness perception. Intrinsic image models are those based on the idea that the visual system segments the image into its illumination and reflective layers. In my chapter I identified such models with an 'illumination-interpretative' stage of lightness perception, whose purpose was to discount the effects of spatially varying illumination (shadows, highlights, shading, transparency) in order to achieve lightness constancy.

Based on my memory of Alan's talk, I have illustrated his demonstration in Figure 3. We were asked to compare the lightnesses of two square patches on a single sheet of black card. The left side of the card was brightly illuminated by a theatre lamp from below. The unoccluded part of the lamp's beam fell on the wall some way above the card, such that the background against which the card was seen was illuminated only by room light, i.e. not the lamp. As Alan pointed out, it should have been obvious to anyone that the left half of the card was bathed in a very bright light, since both the lamp and its shadow were clearly visible, and a penumbra could be seen dividing the two sides of the card. The two patches looked more-or-less equal in lightness. When Alan switched the lamp off however, the left patch, now seen in its 'true light', was seen to be much darker than the right patch. According to Alan, this showed that our visual system had failed to take into account the lamp's illumination when estimating the lightness of the left patch. If the visual system had been able to segment the scene into its intrinsic images, i.e. its illumination and reflective components, then we should have correctly perceived the relative lightnesses of the two patches.

Does this demonstration mean that we should reject the idea that intrinsic-image processing is part of lightness perception, as Alan suggests? I suggest not, because the demonstration does not *isolate* the mechanisms involved. Because the brightly illuminated half of the card was seen against a surround bathed only in ambient room illumination, there was a substantial *contrast* between the left half of card and its background, a contrast only present along one side of the right half of the card. In just the same way as selectively illuminating a near-black patch with a very bright light makes it appear white - the classical Gelb effect - the left half of the card and its patch in Alan's demonstration was subject to a partial Gelb effect, and this would be expected to raise the lightnesses of both card and patch. Our visual system may have partially discounted the lamp's illumination when comparing the lightnesses of the two patches, but the effect of any discounting was almost certainly overwhelmed by the effects of contrast. What Alan's demonstration shows is that there are circumstances in which the visual system fails to achieve lightness constancy with respect to spatially varying illumination because other factors, such as contrast, happen to also play a big role in lightness perception.

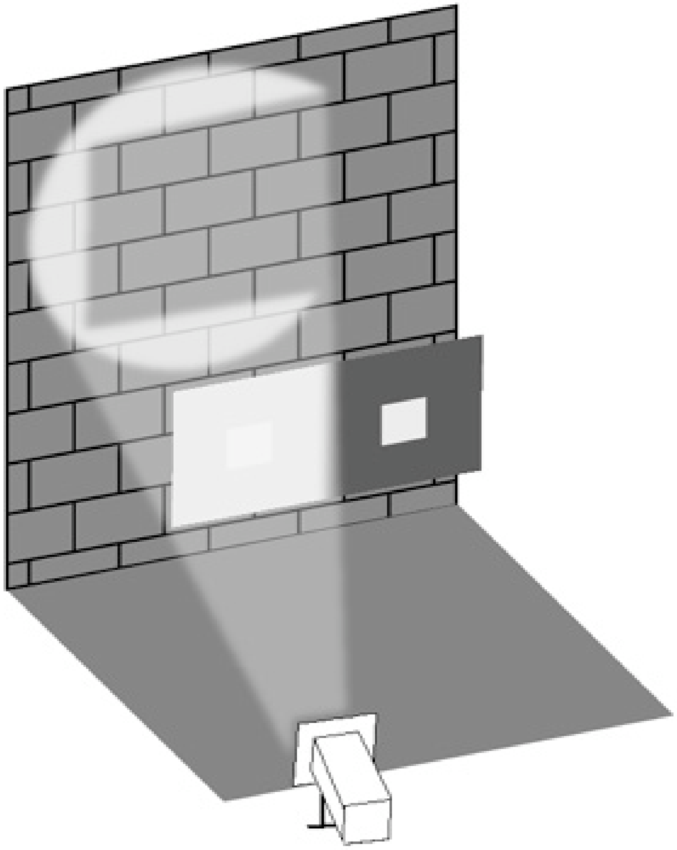

Figure 3. *Demonstration by Gilchrist, from author's memory. The left half of the card was illuminated from below by a powerful lamp, such that its shadow fell on the wall some way above the card, and the card was seen against a background illuminated by the much weaker room light. The two square patches looked more-or-less equal in lightness. When the lamp was switched off however, the left patch was seen to be darker than the one on the right, as illustrated below.*

One cannot help sense an underlying theme in the way Alan uses demonstrations to counter other's models, or rather to counter putative stages of lightness perception. It seems to go something like this. 1. Find a stimulus that demonstrates that a given model/stage is unable to fully explain a given pattern of lightness errors. 2. Reject the model/stage. If one were to apply the same method of falsification to colour vision, our inability to perceive reddish-greens or bluish-yellows would become sufficient grounds for rejecting the trichromatic theory of colour vision, as trichromacy cannot explain colour-opponency.

Does this mean we cannot have critical tests for lightness models/stages ? On the contrary we can, but demonstrations like Figure 3 are not the way. Instead, we need stimuli that *isolate* the potential effects of each putative stage from those others that influence lightness, and in particular contrast. In Figure 13 of my chapter I show an example of a stimulus that I believe achieves this aim. The figure reveals a small but significant effect on lightness of the perceived configuration of illumination relationships, without the confounding effects of contrast. This demonstration thus supports the idea that there is a stage in lightness perception that is 'illumination-interpretative', i.e. that involves intrinsic-image processing.

Absolute versus relative lightness judgements

Alan has often said to me "You cannot have lateral inhibition without anchoring". Well, yes and no. While it is true that a full account of lightness perception must include a mechanism that converts relative lightnesses into absolute ones, by anchoring them to a grey-level standard, I see no contradiction between the need for anchoring and the view that contrast is central to lightness perception. The stock-in-trade of lightness theories are the errors made by the system, such as SLC, and these are errors of *relative not absolute* lightness judgement. In just the same way that I can judge the car in the next lane to be going faster without having to know its actual speed, I am able to judge two surfaces to be different in lightness without having to estimate the absolute shade of grey of either. Contrast theories are able to predict these errors in *relative* lightness judgement without the need for an explicit anchoring stage. As mentioned earlier, Alan believes that anchoring is itself responsible for errors such as SLC, because it operates not only globally but locally within different perceptual frameworks. This is an interesting idea worth pursuing. However, there is no *logical* reason why anchoring might not instead be a stage in lightness perception that takes as its input a map of relative lightness values provided by the combined operation of low-level contrast-sensitive and mid-level illumination-interpretative mechanisms. In this scenario of course, anchoring would not be responsible for errors of relative lightness judgement such as in SLC.