

This article was downloaded by: [Canadian Research Knowledge Network]

On: 3 May 2011

Access details: Access Details: [subscription number 932223628]

Publisher Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Australasian Journal of Philosophy

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713659165>

Reflectance realism and colour constancy: What would count as scientific evidence for hilbert's ontology of colour?

M. Chirimuuta^a

^a Monash University,

To cite this Article Chirimuuta, M.(2008) 'Reflectance realism and colour constancy: What would count as scientific evidence for hilbert's ontology of colour?', Australasian Journal of Philosophy, 86: 4, 563 – 582

To link to this Article: DOI: 10.1080/00048400802215364

URL: <http://dx.doi.org/10.1080/00048400802215364>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

REFLECTANCE REALISM AND COLOUR CONSTANCY: WHAT WOULD COUNT AS SCIENTIFIC EVIDENCE FOR HILBERT'S ONTOLOGY OF COLOUR?

M. Chirimuuta

Reflectance realism is an important position in the philosophy of colour. This paper is an examination of David R. Hilbert's case for there being scientific support for the theory. The specific point in question is whether colour science has shown that reflectance is recovered by the human visual system. Following a discussion of possible counter-evidence in the recent scientific literature, I make the argument that conflicting interpretations of the data on reflectance recovery are informed by different theoretical assumptions about the nature of colour, and of perception. If this is so, there cannot be neutral empirical evidence on this point, and this does seem to undermine Hilbert's claim for empirical support. In the end, I suggest alternative ways of thinking about the relationship between colour ontology and empirical work on colour.

I. Introduction – Hilbert's Case for Scientific Support

David Hilbert first made the case in the nineteen eighties that colours could be identified with *spectral surface reflectances*,¹ and this *reflectance realism* went on to become an influential theory in the philosophy of colour. What is more, Hilbert's treatment of the then new field of computational colour vision set a standard for philosophers' engagement with science. The old debate between realists and antirealists over whether or not colours are objective physical properties has carried on in the last two decades, but with ever more reference to current science—it no longer appears acceptable that philosophers of colour should be unaware of the scientific state of the art. A major selling point of any philosophy of colour is, now, that it be consistent with or even supported by colour science, and Hilbert's theory was one of the first to put the science to centre stage in this way.

My aim in this paper is not to give grounds for a refutation of reflectance realism; rather, my focus on Hilbert's case for scientific support is a means to explore the relationship between colour science and philosophy. What I

¹Scientific terms will be discussed more fully in section II. But briefly, *reflectance* is a property which indicates how much of light incident on a surface is reflected rather than absorbed. *Surface spectral reflectance* (SSR) measures how the proportion of light reflected varies with wavelength of the incident light. *Reflectance recovery* is the putative process of inferring reflectance properties of surfaces from ambiguous photoreceptor signals which confound information about incident light and reflectances.

will be taking issue with is Hilbert's assumptions, first, of *consensus* in colour science, and secondly, of the *neutrality* of the science with respect to the philosophical issue. I will be asking if there is actually any way for science to give an answer to the question of reflectance recovery that is neutral with respect to theoretical commitments about the nature of perception. If this is not possible, it may mean that the lack of scientific consensus that I will be reporting is not due to insufficient data, as it may at first seem, but to underlying conceptual differences. What is more, if the science does not turn out to be innocent of philosophical assumptions, it may be that philosophers' appeals to the science are in some sense circular.

To begin, I note that Hilbert does not cite scientific evidence for colour realism itself. Rather, he argues from folk theory, observing that, 'Pre-reflective common sense is robustly realist about colours . . . ' [Hilbert 1987: 2] and like Lewis [1997] and Jackson [1998: 87–112], Hilbert backs colour realism in order to champion common sense. His project is to select from physics the property that best fits the colour job description set out by folk theory. So it is here that the science comes in: even though Hilbert takes it that the debate between realists and anti-realists is a conceptual one [Hilbert 1987: 16], any particular version of realism must stand up to the tribunal of empirical science. First, the philosophical theory must not have factual implications that are in conflict with the known data. Secondly—a point which is more central to Hilbert's use of the science, and the main topic of this paper—a realist must not identify colours with a physical property which is perceptually unavailable:

I will argue that colour is identical with the characteristic ways objects have of reflecting light The point here is that it could turn out that it is not possible for our perceptions of colour to be correlated with the proposed candidate for colour. If it is impossible for the visual system to determine whether or not an object has the proposed property, then colour, in such an analysis, would turn out to be epistemologically inaccessible. There would be very little point to defending an objectivist view of colour that has as a consequence that we are never able to see the colours of things.

[Hilbert 1987: 17]

That is, Hilbert states that his theory is falsifiable by empirical science in this crucial way, that the theory falls if reflectances are not recoverable by the human visual system. Hilbert then goes on to cite the evidence that this is not the case, writing of the more recent computational work of Maloney and Wandell that:

The success of a theory such as theirs also supports the identification of colour with reflectance. Their work applies an algorithm designed to recover information about reflectance from light to the explanation of features of the psychology of human colour vision. In so far as it is successful, it provides evidence for the view that when we learn the colour of a thing we have learned something about its reflectance.

[Ibid.: 128]

Hilbert's position with respect to scientific support appears unchanged in the more recent statement of reflectance theory [Byrne & Hilbert 2003: 9], so it is interesting to ask if now, after twenty years, the scientific evidence is still in Hilbert's favour. This is the issue of consensus which is to be discussed in section II—*has there arisen any conflicting evidence, that reflectance is not recovered?* But note that Hilbert is assuming not only that the scientific consensus is in his favour, but also that this science is philosophically neutral, that the question of reflectance recovery is straightforwardly empirical.² I will present reasons for challenging this assumption in sections III and IV, and to conclude I will suggest how these observations point to different ways of thinking about the relationship between the science and philosophy of colour.

II. Possible Counter-Evidence on Reflectance Recovery

Surface spectral reflectance (SSR) is a continuous function of the proportion of light reflected by the surface versus wavelength of light. It can be estimated accurately by a spectrophotometer, a physical instrument which samples reflected light at very many different wavelengths. Yet the trichromatic human visual system only has three broadly tuned cone photoreceptor types, which preferentially respond to different portions of the spectrum of visible light. Such undersampling of the spectrum has suggested to some that human estimations of reflectances could only be so crude and inaccurate as to mean that we cannot, in effect, recover reflectance.³ However, the deeper problem is that the spectrum of incident light arriving from a surface is the product of the spectrum of the illuminant as well as the surface's reflectance spectrum. This means that even a high-resolution instrument cannot directly measure the reflectance of a surface if it is illuminated by a light of unknown spectrum. One way of overcoming this problem is to first estimate the illuminant chromaticity in order to *discount* its effects and estimate reflectance. This is the 'inverse' or 'reverse optics' strategy, and much of the research on this strategy, which will be discussed presently, has involved the development of algorithms to infer illuminant spectra (see Hurlbert [1998] and Maloney [1999] for important overviews).

²As far as I am aware, Hilbert does not ever discuss the possibility that scientists studying reflectance recovery share his theoretical views. When Hilbert does mention the philosophical commitments of scientists it is to point out the prevalence of subjectivism [Byrne and Hilbert 2003: 3–4], but he does not suggest how the view might influence their empirical work. However, I in no way mean to suggest that Hilbert is a naïve reader of the science, oblivious of any theoretical commitments. The point I am making is that Hilbert's case for scientific support, as put forward in the sections quoted above, does rest on the assumption of neutrality, i.e. the philosophical innocence of Maloney and Wandell's finding that reflectance recovery algorithms can explain features of human colour perception. If I am accusing Hilbert of anything it is of being a selective reader of the scientific literature. There is no charge of naïvety because it is not *obvious* that theoretical commitments are playing a decisive role in the work on reflectance recovery, but this becomes more obvious when one reads the work of other scientists, whose approach is fundamentally different from Maloney and Wandell's. This is what I aim to show in sections III–IV below.

³E.g. Brown [2003: 257], 'Many [computational models of colour constancy] are even designed to reconstruct the full spectral reflectance functions of surfaces; this seems akin to suggesting that the goal of olfaction is to reconstruct three-dimensional molecular models of odorants.'

Now if reflectances, illuminants, or some other feature of the stimulus (see below) are not suitably constrained, the problem of recovering reflectance by inverse optics is *ill-posed*, meaning that given the information available from the photoreceptors there is no unique solution which disentangles surface and illuminant spectra. But exponents of the reflectance-recovery approach to colour constancy, such as Laurence Maloney, have argued that the reflectances of naturally occurring objects are constrained in such a way that reflectance is recoverable—to an approximation⁴—for the human trichromatic system. For example, Maloney [1986] argued from a large data set that naturally occurring reflectances can be accurately described by low dimensional linear models where reflectances are decomposed into a sum of a small number of n linear basis functions. Maloney and Wandell [1986] give a mathematical demonstration that perfect colour constancy can be attained by a trichromatic system if n is less than or equal to two.

As we have seen, Hilbert sees as crucial to the viability of his theory the scientific finding that reflectance is recoverable. Here is another passage from the 1987 book:

The most important result of Maloney and Wandell's work is that there is an algorithm that operating only on the light reflected from a scene recovers the spectral reflectances of the surfaces in a scene. The identification of colour with reflectance does not present the colour vision system with an impossible task.
[Hilbert 1987: 129]

Hilbert goes on to specify a requirement of the science, referring directly to Maloney and Wandell's low-dimensional linear model:

Under certain conditions it is, in fact, possible to visually obtain information regarding the reflectances of the surfaces in a scene. The two most important conditions that must be met are that the number of parameters in the model of reflectances must be smaller than the number of types of photoreceptor and that the number of parameters in the model of lights must be less than or equal to the number of surfaces of distinct reflectance in the scene.
[Ibid.]

Since Maloney [1986], other researchers have gone on to assess the validity of the low dimensional linear model of reflectances. For example, Westland, Shaw and Owens [2000] confirm Maloney's result that more than three basis functions are needed to capture most of the variance in a wide set of reflectances, but do not go on to support his claim that certain characteristics of the human photoreceptors make a three-or-fewer parameter model sufficient. Nascimento and colleagues [2005] and Oxtoby and Foster [2005] present psychophysical data to show that in practice more than five

⁴No vision scientist has argued that reflectance might be more than approximately resolvable by our visual system, nor does Hilbert's theory require this (see Hilbert [1987: chap 6]). Maloney (personal communication) notes that the goal of his work, strictly speaking, has not been to devise algorithms to reconstruct reflectance, but rather to assign invariant colour descriptors to surfaces, and these descriptors are correlated with SSR.

bases are needed in order for natural and artificial images, respectively, to be reproduced from the basis functions in such a way as to be indistinguishable from the original for a human observer.

However, this low-dimensionality restriction on reflectances is just one of the possible constraints which could facilitate reflectance recovery for a trichromatic system and, in fact, Hilbert's theory does not require that the low-dimensional linear model be correct, only that reflectance is recoverable, whatever the constraint happens to be (see Maloney [1999] for the range of possible physical constraints). A trend in recent research has been to use more complex stimuli in colour constancy research than the flat, matte colour 'Mondrians' made popular by Edwin Land [Land and McCann 1971]. A reason for this is that the over-simplified Mondrian world does not contain the cues, such as shadows and highlights, that exist in the real world and might aid the estimation of illuminant chromaticity, that crucial step in the reverse optics strategy. Maloney's research group has been using simulated 3D scenes rendered realistically with powerful computer graphics now available. Thus Maloney, Boyaci and Doerschner [2005] write,

The scenes in the experiments reported here correspond to more complex inverse problems, where accurate estimation of the colour and albedo of surfaces within the scene presupposes that the visual system effectively estimates more about the spatial and spectral distributions of the illuminant. However, these scenes also contained additional candidate cues that specifically provide information about the lighting model. These inverse problems are not ill-posed and we find that human observers seem able to use the illuminant cues that we provide to solve them.⁵

On the other hand, David Brainard, another leading researcher, advocates using real scenes under controlled viewing conditions as colour constancy stimuli because it seems that in the real world there exist some cues used by humans in colour perception of which scientists are still unaware and so are not able to simulate [Brainard et al.: 2003].

So to summarise this brief overview of the literature, it is striking that scientists are yet to reach a consensus on whether or not reflectance is recoverable *in principle*. Some recent analyses and experiments have put doubt on the low-dimensional linear model of reflectances and lights which were originally cited by Hilbert while, on the other hand, many scientists—Maloney included—have moved on to using complex, naturalistic stimuli which offer the possibility of alternative constraints.

A. Lightness Recovery

Yet aside from this 'in principle' debate on reflectance recovery, it is also to be noted that there is ongoing a wider scientific debate about whether the

⁵Note Maloney's explicit mention of colour and albedo (or lightness i.e. achromatic reflectance) constancy phenomena as 'inverse problems'. This is central to his approach to vision, and in section II.A we will see how these theoretical assumptions play out in his work on lightness constancy.

visual system ever uses reverse optics strategies—that is, over whether reflectance would be recovered *in practice*, even if it were theoretically possible. This question has been attacked head-on in some interesting recent work from the laboratories of Laurence Maloney and Qasim Zaidi, another scientist who has worked on colour and whose approach I will be comparing with Maloney's. These experiments do not look at spectral surface reflectance, but at achromatic reflectance, known as *lightness*, which is the overall ratio of incident to reflected light, averaged across all wavelengths. Like SSR, it is a property of surfaces which is invariant with illumination, as opposed to *brightness*, the product of surface lightness and intensity of illumination. Lightness recovery by inverse optics is a straightforward task in comparison to the problem of recovering SSR and is not necessarily ill-posed. So one can examine the question of whether or not the visual system uses inverse-optics procedures in lightness tasks, aside from the question of theoretical possibility which, as we have seen, is as yet unresolved.

This was the issue taken up by Robilotto and Zaidi [2004] in some experiments in which observers were asked to match the lightnesses of uniform grey cups presented across different luminance levels. Figure 1 shows the stimuli used. On each experimental trial, four stimuli would be presented, two at a time at two different luminance levels (figure 1a; note that in the actual experiment, the high (left) and low (right) luminance presentations would be shown sequentially). Three of the four cups would always be made of the same shade of grey paper (i.e. same lightness), while the fourth cup was either of a lighter or darker shade. The observer would be asked to report which of the four cups is made of the different paper. Here, the answer to the question for the stimuli shown in Figure 1a is made obvious in Figure 1b, where all four cups are presented under the same light level. But for the experimental trial, the observer would be required to solve the task by first deciding whether the odd cup appeared in the first or second luminance presentation (a simple discrimination task), and then deciding which of these two remaining cups was different from the cups shown at the other luminance presentation, where the luminance difference makes this tricky.

Robilotto and Zaidi modelled predictions of the results given two possible observer strategies, either an inverse optics strategy or a simple brightness-difference or 'photometer' strategy—picking the one of the four cups which looked most different, regardless of the effect of illumination level. Figure 1c illustrates these model predictions. Threshold, in terms of proportion correct, is plotted against difference in reflectance (lightness) between the odd object and the standard three objects. The odd object may appear in the full luminance presentation (top graphs) or low luminance presentation (bottom graphs). On the left hand side are the predicted results pattern if the observer follows an inverse optics strategy. In this case, the observer's judgment of object correct (the decision as to which of the two cups that differ from each other is also different from the cups on the other luminance presentation) is only limited by the observer's ability to make the prior judgment of the correct side (the simple discrimination task as to which of

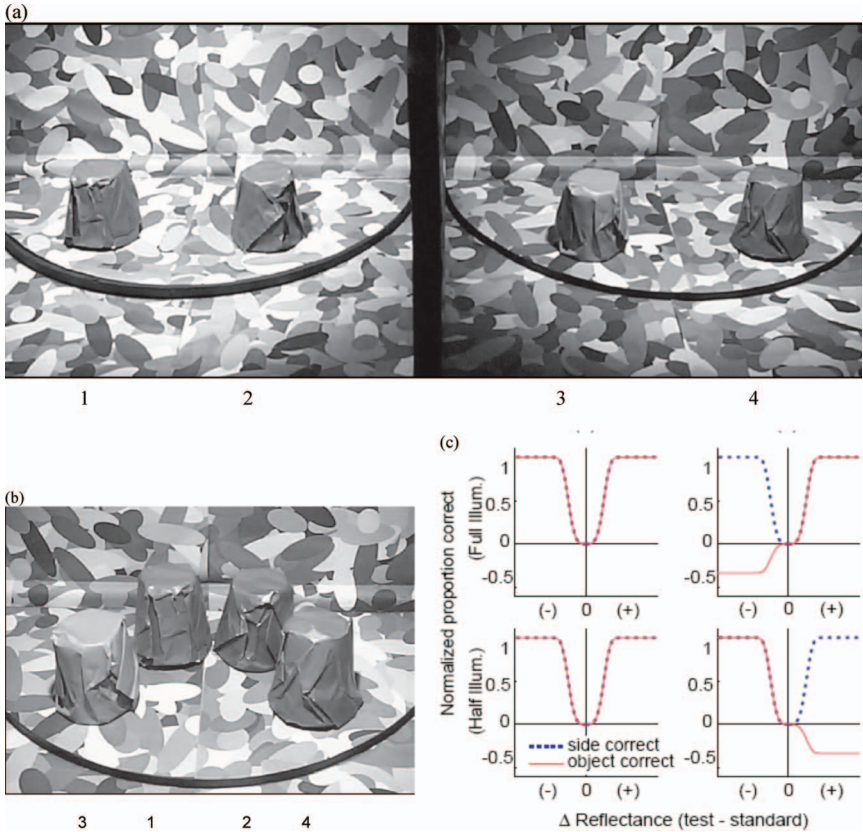


Figure 1. From Robiletto & Zaidi (with permission). (a) Shows the four stimuli as presented, two at a time at two different luminance levels. The observer would be asked to say which of the four cups is covered with a paper that is different (either lighter or darker grey) from that of the other three. The answer is obvious in (b), where all four cups are shown under the same light level. (c) Hypothetical observer strategies for solving the lightness task. Threshold, in terms of proportion correct is plotted against difference in reflectance (lightness) between the odd object and the standard three objects. The continuous line predicts the final decision as to which of the two cups that differ from each other is also different from the cups on the other luminance presentation, and the dashed line predicts performance at the prior discrimination task. The odd-object may appear in the full luminance presentation (top graphs) or low luminance presentation (bottom graphs). On the left hand side is the predicted results pattern if the observer follows an inverse optics strategy, such that lightness judgment is only limited by the threshold for perceiving a difference in appearance of the objects. In contrast, the photometer based strategy (right), in which the observer simply chooses as the odd object the cup which looks most different from the others, will give rise to a characteristic pattern of errors, since the strategy is consistently misleading (the observer will perform below side-correct discrimination threshold and below chance performance, zero on the normalized scale) in certain stimulus conditions—either when the odd object appears in the high luminance presentation and it has a lower reflectance (darker grey) than the standard objects, or if it appears in the low luminance presentation and it has a higher reflectance (lighter grey) than the standards. In these cases (top right graph, left hand side, bottom right graph, right hand side), the object identification (continuous line) pulls away from the discrimination threshold (dashed line).

the two luminance presentations contain the objects that differ from each other). On the right hand side are the predictions for the 'photometer' based strategy, in which the observer simply chooses as the odd object the cup which looks most different from the others in terms of brightness, rather than factoring out the luminance difference to calculate lightness. This strategy will give rise to a distinctive pattern of errors, with object correct and side correct performance pulling apart in some stimulus conditions (see legend).

The data for most subjects were found to be more consistent with the latter brightness-difference strategy, though two observers showed an unusual pattern of errors which suggested that they did factor in overall luminance, *overestimating* the luminance differences (and these anomalies will be discussed below). In a more recent paper, Robilotto and Zaidi present similar results for patterned, as opposed to plain, achromatic stimuli, and they explicitly conclude that their data are evidence *against* the inverse optics model of perception [Robilotto and Zaidi 2006: 33].

Maloney and colleagues, on the other hand, have taken a different approach to lightness constancy and have published results which appear to support the idea that the visual system uses an inverse optics strategy. Boyaci, Maloney and Hersh [2003] looked at lightness-constancy across changes in the angle between target object and light source, realistically rendered in a computerised 3D scene. The perceived brightness of an object changes with angle of illumination, increasing if it is lit directly, whereas the intrinsic property of lightness does not. Therefore if, in a computer-generated scene, the angle of illumination of the light source changes, but the *brightness* of the target object does not, a lightness-constant observer should judge that the *lightness* of the target has changed. This is the pattern of judgment that Maloney and his colleagues tested observers for, comparing observers' results with the predictions of a physical model which solves for perfect lightness constancy and thereby assessing observers' degree of lightness constancy.

Much of the paper is given to a detailed description of this model. The mathematics will not be reproduced here, but a verbal sketch is required for an understanding of the structure of Maloney's argument. The amount of light reflected off a *Lambertian* (matte) surface is a function of the surface albedo (lightness),⁶ the intensity of any diffuse light source, the intensity of a punctate light source, and the angle between the surface and this punctate light source. Boyaci, Maloney and Hersh re-parameterise this model to give light reflected as a function of albedo, total illumination intensity and the 'geometric discounting function', itself a function of the angle of the punctate light source and the ratio of intensities of the diffuse and punctate lights. Such equations, rearranged, would allow one to calculate albedo, given knowledge of the scene lighting and the amount of light reflected from the surface, an inverse optics solution to the problem of lightness constancy.

⁶Maloney and colleagues use the term 'albedo', and Zaidi and colleagues, 'lightness', though the two are equivalent.

On the other hand, knowledge of observers' judgments of albedo allows one to infer the observer's geometric discounting function, that is to find out if the observer is correctly taking into account the intensity and angle of a punctate light source when inferring albedo. As they state:

our goal is to estimate the form of the observer's geometric discounting function and compare it to Equation 3 [the Lambertian, ideal discounting function]. We will allow for the possibility that the observer's perceptions of the layout of the scene, the location of the punctate light source, and the light source intensities are in error. If the observer's geometric discounting function matches the Lambertian geometric discounting function, then the observer is discounting changes in surface orientation in estimating surface albedo.

[Boyaci et al. 2003: 542]

It follows, then, that if observers are not taking changes in orientation into account in their lightness judgments, they cannot be said to be performing an inverse optics calculation.

Figure 2 is the stimulus used. The key object of interest is the central grey rectangle which is the test patch. On each trial, the observer was asked to




Figure 2. The stimulus used in the experiments of Boyaci, Maloney and Hersch [2003] (with permission). The observer must match one patch from the column of reference shades on the right to the central test plane (light grey central rectangle). The other objects in the realistically rendered scene give cues to the angle of spotlight illumination which, Maloney and colleagues propose, the observer may use to infer the intrinsic lightness of the patch from its apparent brightness. This would be an inverse optics strategy. [N.B. To see these diagrams in colour, go to the article's online version.]

judge the albedo of this test plane by matching it to one of the patches from the column of reference shades shown to the right. The other objects in the scene generate highlights and shadows, clues which may be used to estimate the intensity of a punctate light source (a spotlight) relative to ambient illumination, and its angle relative to the test patch.

Boyaci, Maloney and Hersh found, in contrast to some earlier psychophysical results, that observers do take orientation into account in their judgments of lightness, but that constancy performance fell about midway between no-constancy and the theoretical ideal. They used the observers' discounting function to infer what values of the angle and the intensity of the spotlight each observer would have (implicitly) estimated to do the task, on the assumption that the observer is performing the Lambertian model computation. It was found that observers consistently underestimated the intensity of the spotlight relative to the diffuse light, but that they were fairly accurate in their orientation judgments. The latter finding, they conclude, 'suggests that the visual system is effectively estimating information about the spatial organization of the illuminant and using it to arrive at estimates of surface albedo' [ibid. 553]. That is, that the observer is performing an inverse optics computation, first estimating properties of the illuminant in order to discount the illuminant and recover lightness.

In sum, the Maloney group's argument is an appeal to the explanatory success of the model: if one assumes that the visual system is performing a computation equivalent to the Lambertian model, but inaccurately estimating certain parameters, they argue, one is able to account for the constancy data. In another paper, Boyaci, Doeschner and Maloney [2004: 666] compare their results in a colour constancy task with those of Brainard's [1998] who, they write, 'finds that observers' deviations from colour constancy can be parsimoniously explained by the assumption that they have misestimated the chromaticity of the illuminant.'

More can be said about the nature of the scientific evidence and argument for and against lightness constancy since, as the reader will have noticed, there is a significant difference between Maloney and Zaidi's lines of arguments. This difference is due largely to the fact that they are taking positive and negative stances, respectively, on the reflectance recovery hypothesis. While Zaidi argues that his data are inconsistent with the hypothesis (and more consistent with an alternative), Maloney shows us that one way in which the hypothesis would be falsified—if observers were shown to have no lightness constancy—is not suggested by his data, and that an inverse model can give a good account of these data. Note also that Zaidi and Maloney's groups are using different lightness constancy paradigms which, arguably, amounts to an exploration of different problems. In the Zaidi case, constancy is tested with respect to changes in intensity of diffuse illumination; in Maloney's, against changes of angle of a punctate illuminant. Crucially, one group's model cannot be tested with the other group's data. With the current state of the evidence, the two models—and therefore, the two stances with respect to lightness recovery—cannot be compared directly against each other.

So, to conclude this section, again we can note a lack of consensus amongst scientists over reflectance recovery. It might also be pointed out that the case against reflectance recovery is not a 'non-computational' one (in a wide sense), since both Maloney's and Zaidi's studies applied similar methods, psychophysical experiment with computational modelling of the data. However, there is a difference in theoretical or conceptual approach between these scientists which could be called 'Marrian' and 'Non-Marrian'.⁷ I will be arguing that it is important to bear this difference in theoretical perspective in mind when interpreting the significance of this debate. The difference will be the subject of section III, and it will lead us to the question of philosophical neutrality of the science, the subject of section IV.

III. Traditions in Vision Research, Marrian and Otherwise

In section II we saw that recent research has challenged the scientific findings that Hilbert in 1987 cited as support for reflectance realism. On the one hand, the claim that surface spectral reflectance is recoverable by the trichromatic human visual system has long been contested. On the other hand, some researchers have tried to show experimentally that the visual system does not recover achromatic reflectance (lightness), even when this would be a theoretical possibility. However, none of these counter-claims has been conclusive—a consensus is yet to emerge over the scientific questions crucial to the viability of Hilbert's reflectance realism.

What should philosophers of colour do, given the lack of a scientific consensus? One option would be to take up the examination of the philosophical arguments for and against reflectance realism while waiting for the empirical evidence to give its final verdict. Yet this is to assume that a consensus will emerge sooner rather than later, once new experiments are under way and new data are published. In section IV I will be telling a story about why this may not be so, and why this undermines the assumption that empirical evidence about reflectance recovery is philosophically (i.e. conceptually) neutral. As a preliminary, though, in this section I will discuss the conceptual underpinnings of these contrasting scientific approaches to reflectance.

It should now be noted that the scientists who have presented data that suggest that reflectance is recovered, those in the Maloney group, are very much associated with the 'Marrian' school of vision research. This approach to vision was developed in the MIT artificial intelligence laboratories in the 1960s and 70s by David Marr, B.K.P. Horn, and Tomaso Poggio, amongst others. As mentioned above, an inverse or reverse optics strategy is an inference from some sort of sensory signal to a judgment about what physical property or feature gave rise to that signal. In his book *Vision*, Marr describes vision as the solution to a series of ill-posed inverse problems, such as the construction of 3-D stereoscopic information from 2-D retinal arrays, and the example of

⁷Some readers may also be reminded of the issue between Descartes's and Berkeley's theories of vision. See, e.g., Atherton 1990. I would agree that the conceptual difference between Maloney and Zaidi shares something with this longstanding philosophical debate, but I do not trace out the links in what follows.

colour is also taken as a paradigm case [Marr 1982: 17]. So, the tradition is characterised by a conception of vision as a problem of inverse optics.

David Brainard, another leading researcher in the area of computational colour constancy, whose work on ‘equivalent lighting models’ is close to that of Maloney’s (see especially Bloj et al. [2004] and Ripamonti et al. [2004]), has stated very clearly how the Marrian approach can be applied to colour constancy:

[David Marr] articulates the view that vision can be understood as a system that extracts an explicit representation of the world from the retinal image, and that our understanding of human vision is usefully informed by consideration of machine vision algorithms that accomplish the same task. . . . I viewed it [colour constancy] as a relatively simple model problem that embodies the general processing task faced by vision: how can the visual system create a useful representation of surface properties (e.g. colour appearance) from a retinal image that confounds the physical properties of surfaces with those of the illuminant?

[Brainard et al. 2003: 305]

As we have seen, Maloney’s work on colour constancy has framed the problem in exactly this way, and given it one of the most rigorous treatments. Maloney himself has written that,

Visual systems with colour constancy have an objective capability: They remotely sense surface spectral information and represent it through colour.

[Maloney 1999: 389]

What is striking is that this statement is completely congruent with the theory of colour realism or objectivism: Maloney is not saying, like Hilbert, that colours *are* reflectances, but he does say that what our colour vision is about—to the extent that we have colour constancy—is the sensing of reflectances and the labelling of these as colours.⁸

On the other hand, Zaidi has expressed a different set of theoretical opinions which, we will see, do bear on his empirical work on colour. In a 1998 paper he applies a *heuristics* approach to the problem of colour constancy. Quoting Forsyth [1990], Zaidi takes the view that ‘it is neither correct nor helpful to see colour constancy as a problem of measuring surface reflectance’ [Zaidi 1998: 1770]. Rather than recover surface reflectances, Zaidi’s model is designed to match surfaces appearing under different illuminants, for if a visual system can tell which surface colours are the same, despite a change in illumination, it can be said to have relative colour constancy.⁹ Zaidi’s algorithm orders the responses under different illuminants of colour opponent neural pathways or channels, the downstream recipients of photoreceptor signals. This processing shows which

⁸However, Maloney (personal communication) has emphasized that an important difference between his own understanding of inverse problems and Hilbert’s reflectance realism is that Maloney sees no requirement for explicit representation of recovered properties (reflectances), or of any of the estimates of illuminant chromaticities used to recover them. Though this may not be a requirement of Hilbert’s either – see note 10.

⁹Foster [2003] distinguishes this from colour constancy in a more absolute sense, which would entail some sort of reflectance recovery.

responses would be due to the same surfaces, given the different illuminant, and no representation of reflectance is required or aimed at.¹⁰ The ordering of the responses is made possible by certain regularities in the way in which typically occurring luminance changes will effect the relative responses of the colour channels. These regularities make feasible what Smithson [2005] calls the 'RGB heuristic', a way of simplifying the colour constancy problem which is in some ways equivalent to the simplification assumed by the linear models of reflectance.

Cornelissen, Brenner and Smeets give a concise statement of the heuristics approach to colour vision:

The colours we perceive are the outcome of an attempt to meaningfully order the spectral information from the environment. These colours are not the result of a straightforward mapping of a physical property to a sensation . . .
[Cornelissen et al. 2003: 26]

Given the difficulty of disambiguating illuminant and surface information, they frame the problem of colour constancy in the following way:

For spectral information to be useful, one must be able to distinguish surface properties from those of the illumination. . . . Humans and many other animals can somehow recognize colours under a wide range of illuminations. . . . That they are able to do so, can be attributed to the ingenuity of their colour vision systems, which, in many ways, can be understood to be a collection of "tricks."
[Ibid.]

Again, it is to be emphasised that this alternative heuristics approach is equally computational in that it involves mathematical analysis of the problem of colour constancy and the development of algorithms that can match and predict human constancy.¹¹ The contrast with the Marrian tradition is over the conception of the nature of vision and therefore the nature of colour constancy: the Marrian tradition sees the goal of vision as the achievement of veridical representations of the external world by inverse

¹⁰One objection that might be raised here by a proponent of reflectance recovery is, 'What would be matched across illuminants if not reflectance? Reflectance realism might not, in the end, require any explicit representation of reflectance, only correlation between perceived colour and SSR.' But what is different about Zaidi's approach might be brought out with the analogy with light adaptation: when, because of retinal light adaptation, one sees a white sheet as near equally bright inside as outside in the sunshine, one might say that the visual system has matched a property of the object (namely, lightness or albedo). That seems an artificial way of putting it, though. Light adaptation stops photoreceptor responses carrying on increasing linearly to saturation as overall illumination levels increase. This is useful because it means that retinal responses do not saturate out, and that the magnitude of responses to specific objects can stay approximately the same as the illumination intensity goes up. A more natural way of describing the mechanism is that adaptation aims to keep the retinal response to one environmental feature (overall illumination) steady so that differences in another environmental feature (object lightness) will always be detectable, being within the dynamic range of the cells' responses. One might say that lightness is being recovered, but that is an unnecessary step. My point is that the heuristics approach to colour constancy is like this, explaining visual function as a trade-off between sensitivity to environmental changes that are more interesting or less interesting to the animal.

¹¹The difference between the heuristics and the Marrian approaches is subtle because the two views are not mutually exclusive in principle: it could turn out that the heuristic that best captures human colour performance is an inverse optics algorithm. But the heuristics school is not committed to inverse optics in the way that Marrians are. What is crucially different is that the two stances recommend diverging lines of experiment—they constitute different 'research programmes'.

optics, whereas the conception articulated by Zaidi takes vision to be the processing of sensory information in such a way as to find the set of relations between ‘sensory qualities’ (see start of section IV below) that is most useful to the organism.

The heuristics approach is not mentioned by Hilbert in the 1987 book, nor later by Byrne and Hilbert [2003]. It may be worth speculating on why it is that Hilbert does not discuss this alternative approach to colour constancy. One possible reason is that at the time of the first development of reflectance realism (during the mid-eighties), the linear models approach did appear to be promising a completely satisfactory solution to the problem of colour constancy, that is, an algorithm that would predict human constancy performance in the real world, and that could replicate human colour vision in computers or robots. Perhaps this promise has not been fulfilled by linear modelling, although the Marrian inverse-optics program in colour constancy is still very active and productive. Furthermore, twenty years ago an alternative school was yet to be seen to flourish. In the nineties, the availability of greater computing power made possible complex analyses of real-world spectral scenes and so with this new information about natural scenes, and corresponding analyses of how the early visual system responds to them (see e.g. Nascimento et al. [2002]), it became possible to test directly the assumptions that would validate much-mooted heuristics such as the von Kries [1878] or Ives [1912] adaptation formulae (see Smithson [2005] for review).

To finish this section, I would like to emphasize that there is a striking consonance between the conceptual outlook of scientists working in the ‘Marrian’ tradition, and Hilbert’s philosophical position. For example, Maloney has written:

when we study how well human observers judge properties of the environment, including shape, or depth, and separation (length), we usually know, or can determine, the correct answer to any question that we pose to the observer ... we have agreed-upon measurement procedures ... (e.g. a ruler) ... For colour perception, we typically don’t know what counts as the right answer. We don’t have measuring devices to tell us the (true) colour of an object. ... The first contribution of theory to the study of colour perception, then, is development of explicit models of what might count as the physical properties corresponding to colour. Implicit in the structure of such a theory is a claim that there is no fundamental difference between colour, on the one hand, and length or shape, on the other.

[Maloney 2003: 329]

Here Maloney is expressing a version of philosophical colour realism, and moreover, he is saying that colour science could be well served by a theory exactly like Hilbert’s reflectance realism. In particular, he states the working assumption that colours are primary qualities like the others, which is one of the hallmarks of realist colour ontologies.¹² If Maloney and Hilbert agree on the philosophy, perhaps it is should not be surprising that Maloney’s work

¹²E.g. Jackson [1998] calls his microphysical realism ‘The Primary Quality View of Colour’.

supports Hilbert's. But that is to suggest that Maloney's empirical evidence is not impartial. These issues of scientific consensus and neutrality are the subject of the next section.

IV. Consensus and Neutrality

Section II above reported on a lack of consensus amongst scientists now working on reflectance recovery. Section III discussed the difference in conceptual background of two scientists whose opinions differ over reflectance. How are these two matters related? One way to account for any lack of scientific consensus is to say that different scientists have produced conflicting data, and that with refined experimental techniques the difference will be accounted for and consensus will emerge. Yet, I will argue, there is reason to think that in the case of reflectance, the difference in opinion stems from a difference in the conceptualization of the problem of vision, and so does not boil down to a disagreement that can be settled, in any straightforward way, with new experimental output. And this, I believe, is a challenge for Hilbert's idea of scientific support for reflectance realism; for if the scientific opinion in favour of reflectance recovery is more of a theoretical commitment than an empirical discovery, and the theoretical commitment implicit in the science is no less than a version of reflectance realism, the claim of scientific support begins to look circular.

Before racing ahead to this conclusion, however, the case that the difference of scientific opinion about reflectance recovery is as much a conceptual impasse as a conflict over facts needs first to be made. One might summarize the conceptual difference between the Maloney and Zaidi groups by saying that what is at issue is really the *definition of perception*. And furthermore, that this difference radically influences their ways of interpreting fairly similar psychophysical data. We saw in section III that the Maloney group assumes a Marrian, inverse optics, account of vision. On this way of conceptualizing vision, it makes sense to look for a physical property (distal stimulus) that it is the goal of the visual system to represent (to an approximation).

Robilotto and Zaidi, on the other hand, give a different account of perception. They make the distinction between 'sensory' and 'non-sensory' qualities. This is related to the traditional distinction (originally Thomas Reid's) between sensation and perception. A standard textbook [Coren et al. 2004: 8] explains sensation in the following way:

The study of sensation, or sensory processes, is concerned with the first contact between the organism and the environment. Thus, someone studying sensation might look at the way in which electromagnetic radiation ... is registered by the eye. This investigator would look at the physical structure of the sense organ and would attempt to establish how sensory experiences are related to physical stimulation and physiological functioning. These types of studies tend to focus on less complex ... aspects of our conscious experience. For instance ... how we perceive brightness.

While the distinction with perception may be made like this,

[T]he sensory question might be “How bright does the target appear to be?” whereas the perceptual questions would be “Can you identify that object?” “Where is it?” “How far away is it?” and “How large is it?” In a more global sense, those who study perception are interested in how we form a conscious representation of the outside environment and in the accuracy of that representation.

The authors go on to point out that the distinction is not drawn universally by researchers: ‘some investigators have championed its use, and others have totally ignored the difference, choosing to treat sensation and perception as a unitary problem.’

So, on my analysis, Zaidi and colleagues are ‘championing its use’, and Maloney and colleagues are ignoring the distinction.¹³ Therefore, in the context of the lightness experiment, Robilotto and Zaidi introduce the categorization of ‘sensory’ versus ‘non-sensory qualities’, positing that brightness is a ‘sensory quality’ that can be accounted for by early physiological mechanisms, whereas lightness is a ‘non-sensory quality’, only knowable by inference from sensory qualities, which is a perceptual or even cognitive process. In diverging from the Marrian tradition and its principle that vision aims at recovery of non-sensory qualities, Robilotto and Zaidi focus their investigation on the way in which the visual system uses heuristics to solve particular tasks given the sensory cues available. That is, Robilotto and Zaidi are not committed to the idea that vision just is the solution of inverse problems, and so remain agnostic about whether the visual system does go on to recover distal or object properties such as reflectance (‘The visual system may have evolved to identify object properties, but this identification can only be based on sensory information’ [Robilotto and Zaidi 2004: 793]). Their investigation aims to find the sensory (or ‘proximal’) quality that accounts for (i.e. sets the threshold for) psychophysical performance:

In colour matching, an observer does not match spectra, but rather the outputs of cones. We wanted to find out the proximal quality that is used in lightness identification of surfaces. We suggest that, for our 3D objects, this quality is perceived brightness.

[Loc. cit.]

Now, Zaidi’s and Maloney’s groups present somewhat similar psychophysical findings in that both demonstrate that observers have partial lightness constancy. (That is, as with colour constancy, judgment of the lightness of a surface is to some extent stable with respect to illumination level changes, but the stability breaks down in certain situations.) Yet the two groups give

¹³In this section my intention is to remain neutral as to whether the distinction between sensation and perception is justifiable and useful. It might be objected that brightness appearance is not a ‘sensation’ because the psychophysical evidence shows that it is the result of much visual processing. But my point is just that Zaidi and colleagues treat brightness as a sensation (a proximal stimulus), comparing it to lightness which would be a perceived distal property. Nor am I suggesting that, in ignoring the distinction, the Maloney group is denying sequential processing in the visual system.

radically different interpretations of their findings, one group arguing in favour of reflectance recovery, the other against it. How can this be? Robilotto and Zaidi do in fact discuss Maloney's findings, and the key details to bear in mind are the individual differences between observers, including Robilotto and Zaidi's two 'odd' observers, mentioned above, who showed data more consistent with lightness recovery. Robilotto and Zaidi note that, in the Maloney group's work, 'individual differences have been modeled in terms of different estimates for the ambient illumination' [Robilotto and Zaidi 2004: 792]. However, they suggest that these differences are 'likely to be due to attempts to infer a non-sensory quality, rather than due to the particular task or instruction' [loc. cit.]. The crucial point is that the investigation of vision for Robilotto and Zaidi *is* the study of sensory qualities; vision, on their definition, is not the Marrian inference from sensory qualities to object properties. They propose that the Maloney group's reports of illumination estimations are conscious attempts on the part of the observers to infer 'non-sensory qualities', and so are more of a cognitive-perceptual judgment than a visual-sensory process. These attempts, they argue, lead to idiosyncratic differences between observers, and this is also how they account for Boyaci, Maloney and Hersh's [2003] data and their own two odd observers' results.

What is important here is that Robilotto and Zaidi's definition of vision influences the interpretation of the data in such a way that what is a crucial stage in vision in the Maloney interpretation (illumination estimation), becomes a side issue for Robilotto and Zaidi, not really part of the visual process even if it does take place. The upshot is that there are differences in the conceptualization of vision in play which effectively amount to changing the subject or explanandum. If what is at issue is a conceptual difference of this sort, there is little reason to hope that recourse to just more data will be able to settle the matter.¹⁴

Now where does all this leave Hilbert's notion of empirical support for reflectance recovery? The situation might be summarized like this: scientists have theoretical commitments, which may be more or less explicit. Fortunately, for the purposes of this paper, the scientists who have been debating the issue of reflectance recovery are fairly explicit about the conceptual frameworks that they are applying to problems in vision. Interestingly, the theoretical commitments of a scientist such as Maloney map smoothly onto a worked-out philosophical doctrine such as reflectance realism. That is, it can be said that the Marrian framework *implies* colour realism, by positing that, since vision *is* the recovery and representation of external world properties, there must *be* a physical property that we see as colour. And this is entirely consistent with Maloney's explicit statements about the objectivity of colour.

With Zaidi, on the other hand, it is less clear what thesis in the philosophy of colour would best marry with his conception of vision. I do not think it

¹⁴Talk of 'changing the subject' is reminiscent of the old debate about incommensurability. I do not wish to argue here that the different scientific opinions over reflectance recovery are incommensurable in any loaded sense, though the bearing of these cases on some of the wider issues in the philosophy of science will be touched on in the conclusions below.

should be assumed that his framework is anti-realist. The matter is not to be settled here since what is important for this paper is that Maloney should share theoretical space with Hilbert's philosophy, and that Maloney's verdict on reflectance recovery 'reflects' this theoretical position as much as it does the evidence of facts which, as we have seen, stand open to different interpretation by those with different theoretical commitments. Given the possibility that what looked to be scientific support for reflectance recovery may be more an indication of theoretical commitment on the part of the scientists, and given the existence of scientific alternatives – the case that has been made against reflectance recovery, and the overall lack of scientific consensus—one can only conclude that Hilbert's claim for scientific support is undermined.

V. Conclusions: What *Would* Count?

These final thoughts should move us a little towards an answer to the heading question of this paper, 'What would count as scientific evidence for reflectance realism?' First, I suggest that it is not at all inconceivable that a scientific consensus will emerge one day. Even if the so-far observed phenomena of lightness constancy do not decide between a Maloney-type and a Zaidi-type model, it could well be that eventually one approach will have more success than the other in fitting the data of a range of colour and lightness experiments. In which case, it is likely that the less successful 'research programme' will be dropped and researchers will concentrate their efforts on one approach. So even if, as I have hoped to demonstrate, the question of reflectance recovery cannot be settled straightforwardly by experiment, it is conceivable that at some point the majority of scientists will or will not hold that reflectance is recovered. In other words, the reflectance recovery hypothesis stands or falls with the Marrian project.

If things were to turn in favour of Maloney-type models, then I believe that Hilbert *would* have scientific evidence for his theory. If not, one would be entitled to speak of a scientific refutation of his theory. As Maloney himself has written:

If we eventually conclude that no estimation [i.e. reflectance recovery] theory is an adequate description of human colour perception, then we will likely gain insight into the radical difference between perceptual attributes, such as length, that have agreed upon measurement procedure, and perceptual attributes, such as colour, that do not.

[Maloney 2003: 329]

That is, in the success or failure of his sort of approach, Maloney sees implications for how we should think, on a theoretical or philosophical level, about the nature of colour. And such an observation re-emphasizes and does not contradict what I have said about the lack of neutrality of colour science. For it urges us to think of a theoretical opinion, such as a commitment to a realist ontology, as a presupposition of a line of scientific research which ought to be discarded, *modus tollens*, if the line fails to yield

results. This sort of relation between his theory and colour constancy science is what Hilbert [1987: 17] himself points towards when he writes that,

The work of Maloney and Wandell shows that by assuming that the function of colour vision is the determination of the reflecting properties of objects, it is possible to develop a powerful and coherent theory of colour vision.

But this insight is undermined by Hilbert's immediately then talking of the question of reflectance recovery as if it were an empirical issue ('Their work provides the demonstration of perceptual possibility [i.e. the evidence that reflectance is not perceptually inaccessible] that the anthropocentric realist theory requires.' [loc. cit.]; this is taken as 'striking support' for reflectance realism).

It should come as no surprise to us that scientists have theoretical commitments. There is a tradition in philosophy from Locke to Carnap which says that philosophy (ideally) just *is* the explicit and consistent rendering of these commitments. Hilbert's theory of colour seems to me to be a good example of such under-labouring duty (and I say this not intending to be at all derogatory about the intellectual accomplishment of this work). Given this view, we should not expect science to be philosophically neutral—we might do better to drop neutrality as a desideratum for any piece of scientific evidence that is to be bandied back and forth between philosophers. My point, however, has just been that Hilbert put his case for scientific support in such a way as required neutrality, and this could not be borne out.

What I do believe we still require is at least some minimal sort of consensus – an agreement over fundamental lines of approach in research; and I have argued that since we are yet to see any such agreement over reflectance, all potential evidence for reflectance realism remains contentious. Questions remain over what will have been demonstrated if the science which assumes a piece of theory has success: The world is as the ontology describes it? Are colours real? Are colours reflectances? A good Carnapian would have to say no here, but these are issues for another day.

Acknowledgments

I would like to thank Laurence Maloney and Qasim Zaidi for discussions about their work and comments on the manuscript, and for permission to use the figures. I also thank Ian Gold for reading various drafts, and the AJP referees for their comments. This research was supported by the Australian Research Council.

Monash University

Received: January 2007

References

- Atherton, M. 1990. *Berkeley's Revolution in Vision*, Ithaca: Cornell University Press.
 Bloj, M., C. Ripamonti, K. Mitha, R. Hauck, S. Greenwald, and D. H. Brainard 2004. An Equivalent Illuminant Model for the Effect of Surface Slant on Perceived Lightness, *Journal of Vision*, 4: 735–46.

- Boyaci, H., L. T. Maloney, and S. Hersh 2003. The Effect of Perceived Surface Orientation on Perceived Surface Albedo in Binocularly Viewed Scenes, *Journal of Vision*, 3: 541–53.
- Boyaci, H., K. Doerschner, and L. T. Maloney 2004. Perceived Surface Colour in Binocularly-Viewed Scenes With Two Light Sources Differing In Chromaticity, *Journal of Vision*, 4: 664–79.
- Brainard, D. H. 1998. Color Constancy in the Nearly Natural Image. 2. Achromatic loci, *Journal of the Optical Society of America, A*, 15: 307–25.
- Brainard, D. H., J. M. Kraft, and P. Longère 2003. Colour Constancy: Developing Empirical Tests of Computational Models, in *Colour Perception: Mind and the Physical Worlds*, ed. R. Mausfeld and D. Heyer, Oxford: Oxford University Press.
- Brown, R. O. 2003. Backgrounds and Illuminants: The Yin and Yang of Colour Constancy, in *Colour Perception: Mind and the Physical World*, ed. R. Mausfeld and D. Heyer, Oxford: Oxford University Press.
- Byrne, A., and D. R. Hilbert 2003. Color Realism and Color Science, *Behavioral and Brain Sciences*, 26: 3–64.
- Coren, S., L. H. Wood, and J. T. Enns 2004. *Sensation and Perception*, Hoboken NJ: John Wiley and Sons.
- Cornelissen, F. W., E. Brenner, and J. Smeets 2003. True Colour Only Exists in the Eye of the Observer, *Behavioral and Brain Sciences*, 26: 26–7.
- Forsyth, D. A. 1990. A Novel Algorithm for Colour Constancy, *International Journal of Computer Vision*, 5: 5–36.
- Foster, D. H. 2003. Does Colour Constancy Exist? *Trends in Cognitive Science*, 7: 439–43.
- Hilbert, D. R. 1987. *Color and Color Perception*, Stanford: CSLI.
- Hurlbert, A. C. 1998. Computational Models of Colour Constancy, in *Perceptual Constancy: Why Things Look As They Do*, ed. V. Walsh and J. Kulikowski, Cambridge: Cambridge University Press.
- Ives, H. E. 1912. The Relation Between the Colour of the Illuminant and the Colour of the Illuminated Object, *Trans. Illuminat. Eng. Soc.*, 7: 62–72.
- von Kries, J. 1878. Beitrag zur Physiologie der Gesichtsempfindungen/Physiology of Visual Sensations, in *Sources of Colour Science*, ed. D. L. MacAdam, Cambridge MA: MIT Press. 101–8.
- Jackson, F. 1998. *From Metaphysics to Ethics*, Oxford: Clarendon Press.
- Land, E. H., and J. J. McCann 1971. Lightness and Retinex Theory, *Journal of the Optical Society of America*, 61: 1–11.
- Lewis, D. 1997. Naming the Colours, *Australasian Journal of Philosophy*, 75: 325–42.
- Maloney, L. T. 1986. Evaluation of Linear Models of Surface Spectral Reflectance With Small Numbers of Parameters, *Journal of the Optical Society of America, A*, 3: 1673–83.
- Maloney, L. T. 1999. Physics-based Approaches to Modelling Surface Colour Perception, in *Colour Vision*, ed. Gegenfurtner and Sharpe, Cambridge: Cambridge University Press.
- Maloney, L. T. 2003. Surface Colour Perception and its Environments, in *Colour Perception: Mind and the Physical World*, ed. R. Mausfeld and D. Heyer, Oxford: Oxford University Press. 329–30.
- Maloney, L. T., H. Boyaci, and K. Doerschner 2005. Surface Colour Perception as an Inverse Problem in Biological Vision, *Proceedings of the SPIE - IS & T Electronic Imaging*, 5674: 15–26.
- Maloney, L. T., and B. A. Wandell 1986. Colour Constancy: A Method for Recovering Surface Spectral Reflectance, *Journal of the Optical Society of America, A*, 3: 29–33.
- Marr, D. 1982. *Vision*, San Francisco: W.H. Freeman.
- Nascimento, S. M. C., F. P. Ferreira, and D. H. Foster 2002. Statistics of Spatial Cone-excitation Ratios in Natural Scenes, *Journal of the Optical Society of America A*, 19: 1484–90.
- Nascimento, S. M. C., D. H. Foster, and K. Amano 2005. Psychophysical Estimates of the Number of Spectral-reflectance Basis Functions Needed to Reproduce Natural Scenes, *Journal of the Optical Society of America, A*, 22: 1017–22.
- Oxtoby, E. K., and D. H. Foster 2005. Perceptual Limits on Low-dimensional Models of Munsell Reflectance Spectra, *Perception*, 34: 961–6.
- Ripamonti, C., M. Bloj, R. Hauck, K. Mitha, S. Greenwald, S. I. Maloney, and D. H. Brainard 2004. Measurements of the Effect of Surface Slant on Perceived Lightness, *Journal of Vision*, 4: 747–63.
- Robilotto, R., and Q. Zaidi 2004. Limits of Lightness Identification for Real Objects Under Natural Viewing Conditions, *Journal of Vision*, 4: 779–97.
- Robilotto, R., and Q. Zaidi 2006. Lightness Identification of Patterned Three-dimensional Real Objects, *Journal of Vision*, 6: 18–36.
- Smithson, H. E. 2005. Sensory, Computational and Cognitive Components of Human Colour Constancy, *Philosophical Transactions of the Royal Society B*, 360: 1329–46.
- Westland, S., J. Shaw, and H. Owens 2000. Color Statistics of Natural and Man-made Surfaces, *Sensor Review*, 20: 50–5.
- Zaidi, Q. 1998. Identification of Illuminant and Object Colors: Heuristic-based Algorithms, *Journal of the Optical Society of America A*, 15: 1767–76.