

The Paradox of Colour Constancy: Plotting the Lower Borders of Perception

Will Davies

St Anne's College, University of Oxford

ABSTRACT: This paper resolves a paradox concerning colour constancy. On the one hand, our intuitive, pre-theoretical concept holds that colour constancy involves invariance in the perceived colours of surfaces under changes in illumination. On the other, there is a robust scientific consensus that colour constancy can persist in cerebral achromatopsia, a profound impairment in the ability to perceive colours. The first stage of the solution advocates pluralism about our colour constancy capacities. The second details the close relationship between colour constancy and contrast. The third argues that achromatopsics retain a basic type of colour constancy associated with invariants in contrast processing. The fourth suggests that one person-level, conscious upshot of such processing is the visual awareness of chromatic contrasts 'at' the edges of surfaces, implicating the 'colour for form' perceptual function. This primitive type of constancy sheds new light on our most basic perceptual capacities, which mark the lower borders of representational mind.

The notion of perceptual constancy is receiving renewed interest in the philosophy of perception. Historically, constancy has generated much puzzlement, in how best to account for the dual, competing invariance and variance that we encounter, for example, when looking at a white wall in yellow light, or a coin tilted at an angle. These remain live

concerns, but recently there has been a shift in focus, to theorising constancy as a basis for objective mental representation of our environments. On this theme, Burge (2014: 486) has argued that ‘the most primitive, distinctively psychological type of representation is perception. The lower border of perception is the lower border of representational mind.’ In turn, the borders of perception are marked by the presence of perceptual constancies. Constancies are capacities for ‘objectification,’ for ‘marking off states that are as of specific system-independent elements in the environment from states idiosyncratic or local to the perceiver,’ (2010: 400). This ‘notion of objectivity provides the most specific constraint on the type of sensory transactions that count as perceptual,’ (2010: 396).

Understood in this way, constancies provide focal points for enquiry into the very nature of perception. Correlatively, specific perceptual capacities, such as visual shape, orientation, or colour perception, may be approached by unpacking the details of constancies for shape, orientation, or colour. We can distinguish two ways of tackling this project. The first, which is characteristic of much philosophy of perception, focuses on the complex, fully-fledged constancy capacities involved in our (presumably) normal adult human experience. The second, less frequently encountered approach targets the other end of the continuum: what are the most primitive types of perceptual constancy? What are the basic constituents or components of the mature capacity? What do these reveal about the lower borders of perception and representational mind? This paper applies this second approach to the case of colour, developing an account of a particularly basic type of colour constancy. So basic, in fact, that it is initially unclear whether it deserves the label ‘colour constancy’ at all.

The account develops as a response to the *paradox of colour constancy* (section 1). The paradox concerns cases of cerebral achromatopsia, a loss of colour consciousness following cortical injury. Complete cerebral achromatopsics lack any awareness of colour, seeing the world in shades of grey. Bizarrely, however, some such patients show signs of

possessing colour constancy. This is inconsistent with the intuitive, apparently tautological view that colour constancy involves constancy of perceived colour. The argument starts by identifying and rejecting a key assumption that helps generate the paradox: monism about colour constancy. I argue instead for pluralism about our constancy capacities, which opens the door for alternative concepts fitted to the achromatopsic case (section 2). I start this constructive project by outlining the fundamental but neglected relationship between colour constancy and contrast. I argue that the preservation of key types of contrast processing justifies the attribution of constancy in achromatopsia (section 3). Although much of this argumentation appeals to low-level facts about the visual system, it clarifies the close ties between two perceptual phenomena – constancy and contrast – which feature in many philosophical debates, but which are often treated quite differently. This argument does, nonetheless, leave open the person-level, conscious bases for this basic constancy capacity. The final section draws up the relevant connections. I argue for the *chromatic edge hypothesis*, on which the constancy in question involves perceptually representing chromatic contrast properties ‘at’ the edges of surfaces. This implicates constancies associated with the so-called ‘colour for form’ function, as distinct from the ‘colour for colouring’ function most canonically associated with colour perception (section 4). Such chromatic edge constancy is one of the most basic, yet elusive constituents of the mature capacity colloquially known as ‘seeing colours.’

1. The Paradox of Colour Constancy

What is colour constancy? There is, of course, substantial disagreement on this issue.¹ I nonetheless think that most parties would accept this first-pass characterisation:

Constancy for Colours: Colour constancy is constancy with respect to the perceived colours of objects viewed under changing/varying illumination.

One finds similar descriptions in the introductory sections of many philosophical and scientific works on colour constancy. To pick a prominent example of each, Hilbert (2005: 141) defines ‘colour constancy’ as ‘constancy with respect to the perceived colours of surfaces’ under changes in lighting, while Arend and Reeves (1986: 1743) stipulate that constancy is ‘a process whereby perceived object colours remain invariant under changes of illuminant colour’. I take it that such characterisations are intended to provide a neutral starting point for theorising.² Prima facie, they seem unobjectionable. Indeed, it seems truisitic that colour constancy is perceptual constancy for colour, which presumably just is constancy with respect to perceived colour.

There are two immediate points on this characterisation. The first is that it is circular, using the term ‘constancy’ on both sides. Arend and Reeves use the term ‘invariant,’ rather than ‘constant,’ but this does not clarify matters. Further explication of this notion is therefore required, but that is only to be expected, given the pre-theoretical status of the definition. The second – and for our purposes, more significant – point is that constancy is defined in terms of ‘perceived colour.’ This point has two important implications. The first is that Constancy

¹ Recent discussion includes Brown (2014), Cohen (2008), Davies (2016, 2018), Gert (2010), Hilbert (2005), and Wright (2013).

² Notably, Constancy for Colours is neutral on the issue of colour ontology.

for Colours entails that any subject with colour constancy must have a capacity to perceive colours. If colour constancy just is constancy with respect to perceived colours, then no perceived colours means no colour constancy. The second, related implication is that we have a prior, independent grasp of the notion of ‘perceived colour,’ such that it may play the envisaged role in helping to explicate the notion of ‘colour constancy.’ That is, ‘perceived colour’ is presumed to be analytically prior to ‘colour constancy.’ What notion of ‘perceived colour’ is in play here? Again bearing in mind its pre-theoretical status, a suitably neutral gloss is that the perceived colour of X (for S , at t) is the colour property that S ’s visual experience at t presents X as having. Plugging this into Constancy for Colours, colour constancy is constancy (invariance) with respect to the colour properties that one’s visual experience presents an object as having, under changes or variations in illumination. And this does seem to be the intuitive view.

Our puzzle concerns putative cases of constancy in subjects who lack the ability to perceive colours in this sense. These are subjects with cerebral achromatopsia (from hereon, *achromatopsia*), a rare condition involving a selective impairment in the ability to have conscious visual experiences of chromatic colour.³ Unlike more common forms of colour blindness, achromatopsia involves cortical damage, rather than retinal abnormalities, (hence, ‘cortical colour blindness’). Mollon and colleagues (1980) found three fully functional retinal cone mechanisms in a complete achromatopsic known as ‘MS’. MS’s spectral sensitivity function was also found to be consistent with the presence of cone-opponent processing (Heywood et al., 1996). Yet due to their brain injury, complete achromatopsics describe their visual world as ‘drab and grey’ (Covey & Heywood, 1995: 90) or ‘drained of colour’ (Heywood et al., 1987: 22). Indeed, they are unable to perform standard diagnostic tests for

³ From hereon, I suppress ‘chromatic.’

colour vision: they cannot discriminate isoluminant stimuli in respect of colour, accurately name the colours, nor produce appropriate orderings of colour chips in the Farnsworth-Munsell 100 hue test.

Despite this profound impairment of colour consciousness, several authors have reported evidence of preserved colour constancy in achromatopsia. A study by Hurlbert and colleagues (1998: 143) concludes that MS has ‘rudimentary colour constancy... mediated by the computation of cone contrasts between image regions.’ Reporting this work prior to publication, Cowey and Heywood (1997: 135) claim that ‘M.S. has normal colour constancy...’ In an influential review article, Smithson (2005: 1342) concurs that ‘a cerebrally achromatic [sic] observer ... [was] able to display colour constancy in an asymmetric matching paradigm.’ Another study by Barbur and colleagues (2004: 25), involving different stimulus conditions, found that ‘either normal or slightly reduced, but functioning, ICC [instantaneous colour constancy] mechanisms are present in [three] achromatopsic subject[s].’

This brings us to the paradox of colour constancy. On the one hand, we have a reasonable scientific consensus that achromatopsics have at least ‘rudimentary’ colour constancy. These reports come from leading experts in the field, and therefore have considerable *prima facie* credibility.⁴ On the other hand, we have Constancy for Colours. This apparent truism tells us that colour constancy involves invariance with respect to the perceived colours of surfaces under changes or variations in illumination. But achromatopsics

⁴ I stress: *prima facie* credibility. The question of whether this capacity in fact warrants the label ‘colour constancy’ is a substantive philosophical issue, which cannot be settled by appeal to empirical authority. Section 2 provides an extended argument in support of classifying the capacity in this way.

cannot perceive colours! The intuitive characterisation of constancy therefore gets no grip here whatsoever.

Clearly, something has to give, but what exactly? The most obvious response is to reject one of the assumptions that generates the paradox. Unfortunately, however, neither option is straightforward. Firstly, we might take issue with Constancy for Colours. In particular, perhaps it was a mistake to define ‘colour constancy’ in terms of ‘perceived colour,’ where this implied a capacity to have conscious visual experiences that present the colours of objects. One might argue that this starting point is insufficiently neutral, in barring colour constancy in human or non-human animals that lack conscious awareness of colour. One alternative would be to unpack ‘perceived colour’ in more functional representational terms, rather than in terms of properties presented in visual experience.⁵ Glossing over important details, such a view might allow for perceptual representations of colour in creatures that lack conscious awareness of colour. On this view, the attribution of constancy in achromatopsia might be vindicated, if these subjects were to possess a non-conscious capacity for representing object colours, which exhibited the relevant type of illumination-invariance. That is, adopting this view may resolve our paradox, if it turned out that achromatopsics possess a kind of blindsight for object colour.⁶

⁵ This approach is discernible in Burge (2010: 368, 374-76, 402) and Matthen (1999: 78).

⁶ ‘Blindsight’ here means *type 1* blindsight, where patients supposedly have no conscious awareness of stimuli presented in the impaired visual field, yet can discriminate some of their properties. *Type 2* blindsighters have some residual awareness of stimuli in their ‘blind’ field, which facilitates discrimination. This residual awareness might be severely degraded or unusual. I consider the possibility that MS has degraded awareness of colour below.

Whatever the merits of the functional representational view, it will not help us here. That is because the putative constancy capacities in achromatopsia, mysteriously enough, seem to involve conscious perception.⁷ I say more on this later, but to be clear, it seems that achromatopsics can perform certain paradigmatic colour constancy tasks using conscious perception, despite apparently lacking any awareness of the colours of the stimuli. No tweaking of Constancy for Colours will help in accommodating this. Only a radical revision will suffice, somehow reshaping the notion to cover constancies that do not even imply the ability to perceptually represent colours. Such an extreme departure is too much to swallow, at least straight off the bat.

Perhaps we should therefore question the scientific consensus. It might have been premature to describe achromatopsics as ‘colour constant,’ purely on the basis of their ability to perform certain psychophysical constancy tasks. After all, it would not be the first time psychologists have over-extended mentalistic notions, in attempting to explain seemingly intelligent or complex behaviours. Perhaps, then, we should look for an alternative, deflationary explanation of the achromatopsic data, which prescind from attributing constancy proper.

Now, although this strategy may seem more promising, it is still problematic. Firstly, we are owed an account as to why the experts should want to attribute constancy in achromatopsia in the first place. After all, it is not as if they are unaware of the puzzling ramifications. Hurlbert and colleagues (1998: 143), for example, note that ‘at first sight, it seems paradoxical to test colour constancy in a cerebrally achromatopsic observer.’ The implication being: at first sight, but not on closer inspection. Secondly, while the previous response seemed worryingly revisionary, the deflationary response arguably seems guilty of

⁷ Heywood et al. (1994: 251; 1998: 413-15).

parochialism. On the face of it, the reports of constancy in achromatopsia look like a discovery, rather than a gross conceptual error. Of course, these claims are hard to understand, given the ordinary concept of constancy. But the history of psychology and philosophy of mind is littered with theoretical advances stemming from previously unimaginable dissociations. Why not here? The possibility at least warrants closer investigation.

As should be clear by now, there is no quick and easy resolution to the paradox of colour constancy. The remainder of the paper takes up the challenge, developing a response that respects, as I think we ought, both Constancy for Colours and the scientific consensus. The first stage of the argument rejects an important, unarticulated assumption that helps generate the puzzle: monism about colour constancy. The resulting pluralism opens the door for alternative concepts, which may depart from Constancy for Colours in fundamental ways – even to the extent of allowing types of constancy in the absence of any perceptual representation of colour.

2. Pluralism

In presenting Constancy for Colours as a definition of ‘colour constancy,’ the implicit suggestion was that the term was monosemous, and that constancy is a singular, unified phenomenon. As a matter of fact, it is not. I have previously defended the following view (Davies, 2016, 2018):

Pluralism: There are different types of colour constancy, involving different perceptual capacities. These capacities are individuated by differences in environmental conditions, computational description, representational content, mechanism, and aspects of phenomenal character.

Pluralism is not motivated by naïve reflection on visual experience, rather by close attention to empirical practice. In vision science, I think that pluralism is in fact the standard view, though this is rarely made explicit. For example, an influential review by Smithson (2005: 1329) details ‘the sensory, computational and cognitive aspects of human colour constancy.’ She distinguishes models based on adaptation, involving multiplicative scaling of cone signals, such as the Ives transform and von Kries coefficient rule; a range of computational models aimed at providing estimates of the illuminant; and models of ‘operational [i.e. relational] colour constancy’ based on ‘coding colour relations by ratios’ of cone-excitation values (2005: 1334). Foster (2011: 696) draws similar distinctions, concluding that,

multiple mechanisms underlie constancy judgments, each providing cues to the state and stability of the observed surface, object, or scene... Which surface-colour attribute is given perceptual prominence may depend simply on the task at hand, but at present it is not possible to identify uniquely either the neural substrate for these attributes or how they are combined with other non-chromatic attributes to determine surface-colour appearance.

Granzier and Valsecchi (2014: 4) echo these thoughts,

colour constancy might be like a “bag of tricks”; the kind of information and the combination of information that will be used by the visual system will depend on the task at hand..., the observer..., and the presence of the information itself....

These remarks speak resoundingly in favour of pluralism.

To illustrate concretely, let’s briefly compare models that estimate the illuminant with models of relational constancy. Simplifying significantly, the computational problem addressed by the former is to produce representations of the colours at each point on a surface that is viewed under changing (or varying) illumination, despite the fact that surface and

illuminant properties are confounded in the proximal stimulation of the visual system. The basic form of the solution is to generate a representation (an ‘estimate’) of the illuminant, in order to ‘discount’ its contribution to the incoming signal. The estimate formation may recruit various ‘cues’ in the scene, such as average chromaticity, specular highlights, shadows, and more besides. It is notoriously controversial as to what the perceptual corollary of such computational processes might be, but the mainstream, broadly empiricist view is that it involves a non-sensory, quasi-cognitive, perhaps conceptualised aspect of perceptual experience; more recently, a ‘projective’ or ‘distal’ mode of perceptual representation.⁸

Relational models, in contrast, involve no estimation of the illuminant. Given some change in proximal stimulation from two points p and q , the problem is to determine whether this was due to an illuminant change, or a change in surface colour properties. The solution, roughly, is to compare the stimulation from p and q before and after the change, with invariant ratios indicating an illuminant change, and variant ratios indicating a surface colour change. (More on this later.) Such relational constancy is strongest under conditions of rapidly changing illumination, and is not significantly attenuated by reducing cues as to the illumination. The perceptual corollary of relational constancy is again up for debate. For now, let us suppose, minimally, that we *in some sense* perceive the colour relations between p and q as invariant under illumination changes, even while their ‘absolute’ colour values may vary quite significantly.

In philosophy, pluralism is endorsed by Wright (2013), and is implicit in Burge (2010).⁹ Wright (2013: 436) argues that there are ‘two forms of colour constancy, one

⁸ For ‘projective’, see Reeves et al. (2008). For ‘distal’, see Palmer (1999).

⁹ Pluralism is also implicit in Brown (2014), who distinguishes between illumination-independent and transparency-independent constancy.

phenomenal and the other projective.’ Following Reeves and colleagues (2008), he associates the ‘phenomenal’ kind with processes of adaptation, while the ‘projective’ kind is tied to processes of ‘inference.’ Wright (2013: 443ff) also discusses relational constancy, and seems to allow that this might be a third type. As will emerge later, my view differs from Wright in positing a greater multiplicity of types, distinguishing illumination-independent and background-independent colour constancy.¹⁰ Burge (2010: 351ff) provides a detailed description of a basic type of lightness constancy, in which the visual system exploits systematic differences in ‘luminance contours in the sensory registration’ of light, in order to distinguish illumination edges from reflectance edges. Roughly, sharp luminance contours typically correspond to reflectance edges, whereas fuzzy contours correspond to illumination edges.¹¹ Burge (2010: 354) adds that ‘the full lightness constancy capacity in humans is more complex than the capacity I have described, which is only a component in the full capacity.’ The notion of a ‘component’ capacity suggests that the edge-categorisation process itself constitutes a type of lightness constancy. This indicates pluralism, though it is also consistent with the monistic view that only the ‘full capacity’ is lightness constancy proper, with the ‘components’ constituting mere parts of the mechanism that sustains it. I shall not rehash the previous points in arguing against the latter view.

Much more could be said about the motivations for pluralism, but this would distract from our main aim: resolving the paradox of colour constancy. To recap, the paradox arose

¹⁰ I also distinguish atmosphere-independent colour constancy, which I will discuss in future work.

¹¹ As Adelson & Pentland (1996) note, however, sharp luminance edges can arise in natural scenes from both reflectance changes and illumination changes. Other heuristics must therefore be at work in distinguishing these distal changes.

by conjoining Constancy for Colours with the scientific consensus concerning achromatopsia. We already rejected the possibility of revising Constancy for Colours to accommodate the putative achromatopsic cases. Such a revision would be too radical, alienating us from any intuitive grip on the phenomenon. Pluralism facilitates a more nuanced approach. Let's begin by noting that, even granting pluralism, all parties should agree that Constancy for Colours is central to our understanding of colour constancy. It plausibly provides an exemplar characterisation, helping to identify core cases of the full mature capacity. But it would be a mistake to define 'colour constancy' with reference to such cases, as requiring 'constancy in respect of perceived colour.' This would preclude – as a matter of definition – a capacity such as relational colour constancy, which involves constancy in respect of perceived colour relations, rather than colours per se. Similarly, defining 'lightness constancy' as 'constancy in respect of perceived lightness' would preclude the capacity that Burge discusses, involving the categorisation of illumination edges and reflectance edges. We should reject such conceptual hegemony. Both capacities involve forming objectifying perceptual representations of distal features: respectively, colour similarity relations between objects, and boundaries of material surfaces and regions of illumination. Both involve the operation of computational principles that function to disentangle the contributions of certain distal features to the proximal stimulation of the visual system. Both capacities constitute perceptual constancies, according to the science. Neither involves 'constancy in respect of perceived colour (lightness),' in the intuitive sense of that phrase.

In tackling our puzzle, then, we may combine a certain conservatism regarding Constancy for Colours, considered as a core conception, with a more open-minded view about the various forms that colour constancy might take, once we unpack the constituent capacities comprising the mature condition. And within this framework, it is legitimate to

ask: does the heralded capacity in achromatopsia in fact constitute a type of colour constancy? If so, what type exactly? The next two sections tackle these questions in turn.

3. Contrast and Constancy in Achromatopsia

In introducing the paradox, the scientific consensus provided *prima facie* grounds for believing that achromatopsics possess a type of colour constancy. Clearly, however, such an appeal to authority provides no cogent philosophical reason for endorsing this attribution. The aim of this section, firstly, is to develop a philosophical framework within which we can start to make sense of this puzzle. And secondly, to deploy this framework to argue that the presaged capacity in achromatopsia does indeed constitute a type of constancy. The framework outlines the close relationship between colour contrast and colour constancy. Both phenomena are, to a significant extent, products of the cone contrast code employed at low levels in the visual system. To substantiate this point, I highlight some properties of cone contrast processing that are especially important to constancy. I then argue that some complete achromatopsics possess the relevant sort of contrast processing, which moreover appears to have upshots in conscious experience, hence should be credited with a very basic form of colour constancy. I refine this claim by considering some objections, and offering replies. The section closes with discussion of the wider philosophical implications of this view.

To begin, we must distinguish two different notions of contrast. The familiar kind is *colour contrast*, which concerns the relationship between the apparent colour of an object, and the apparent colour of its surround. So-called *colour contrast effects* are cases in which the colour that an object visually appears to have differs quite substantially, when viewed against backgrounds of differing colours. Figure 1 illustrates such effects for (a) achromatic and (b) chromatic colours. In 1(a) and 1(b), the two discs have exactly the same surface

properties, but the left-hand disc in 1(a) appears darker than the right-hand disc, and the left-hand disc in 1(b) appears slightly more reddish than the right-hand disc. Such contrast effects have long been a source of scepticism about the mind-independence or objectivity of colour. What colour is the left-hand disc in 1(b) really? It seems that nothing privileges one or other background as more favourable in determining the colour of the disc. As such, there is just as much reason for thinking that the disc has the colour that it appears to have on the right-hand side, as that it has the colour it appears to have on the left. Given the equipollence of these judgements, the only rational options seemingly are either to conclude that the disc has both colours, or that it has neither.¹²



Figure 1

The second type of contrast is *cone contrast*, which concerns the relative activity of receptor cone cells across temporal or spatial intervals. Let's abbreviate the photon catch from point p in cone class L as $L(p)$ – similarly for M or S cones. One simple type of cone contrast is the spatial ratio of cone excitation between two points, within a single cone class: $L(p)/L(q)$. We could also take the temporal ratio of excitation from a single point at

¹² For 'both', see Cohen (2009). For 'neither', see Hardin (1988). Mind-independent realists would reject the initial reasoning. Byrne & Hilbert (2003), for example, argue that the mistake is to assume that if the disc really does have one colour rather than the other, we must be in a position to know which colour it has.

successive times: $L(p)/L(p')$. As will be familiar, the early visual system compares the outputs of different receptor types, yielding at least three opponent channels. The standard, simplified model includes an $(L+M)$ or ‘achromatic’ channel, $(L-M)$ or ‘red-green’ channel, and $(S-(L+M))$ or ‘yellow-blue’ channel.¹³ For our purposes, it will be convenient also to introduce MacLeod and Boynton’s (1979) chromaticity diagram, which represents the two chromatically opponent channels via an $[L/(L+M)]$ or ‘red-green’ axis, and an $[S/(L+M)]$ or ‘yellow-violet’ axis. Let’s represent the $[L/(L+M)]$ value of p as $RG(p)$, and the $[S/(L+M)]$ value as $YV(p)$. The spatial contrast between $RG(p)$ and $RG(q)$ can be represented by the distance $|RG(p)-RG(q)|$. Let’s call these *cone-opponent contrasts*. All incoming spectral information is coded via contrasts of some such sort: the visual system works with relative rather than absolute cone and cone-opponent values. Although this is far from the whole story, contrast signals are of fundamental importance in understanding many colour-related phenomena. As Hardin (1988: 15) augurs, ‘this selective emphasis on contrast at such an early level of visual processing is bound to have a powerful effect on what we see.’

One such phenomenon, predictably, is colour contrast. Consider image 1(b) again. The orange disc presented against the green background signals a greater ‘red-green’ cone-opponent contrast than the disc presented against the red background. As such, a mechanism with combined chromatic and spatial opponency, computing something like $|RG(p)-RG(q)|$ values for the disc/background pairs, would produce higher levels of activity with the orange/green pair than with the orange/red pair. Such ‘double-opponent’ mechanisms abound

¹³ The scare quotes are intended to flag the poorly understood relationship between activity in early cone-opponent channels, and the opponent structure of our perceptual experience of shades of red, green, and so on. It is surely just as much a mistake to label an $(L-M)$ channel a ‘red-green’ channel, as it is to label an L cone a ‘red’ cone.

in primary visual cortex.¹⁴ Now, clearly there is a significant leap from here to facts concerning our visual experience of these pairs. I'm not lining up to deliver the relevant bridging principles. We can confidently say, however, that the manifest differences in our experiences of these pairs are predicted rather well by the differences in spatial cone-opponent contrasts – at least for simple stimuli of this sort.

Another phenomenon that is intricately linked to cone contrast, perhaps less predictably, is colour constancy. To be clear, this connection is common knowledge among vision scientists, having been noted in empirical literature for over two hundred years (Mollon, 2006). Among philosophers, however, as noted above, contrast phenomena are often cited as reasons to think that colour is, in some sense, perceiver-dependent. Conversely, constancy has been touted as providing some of the strongest motivation for mind-independent realism.¹⁵ As we shall now see, however, colour contrast and colour constancy are, to a significant extent, two sides of the same coin, in terms of their underpinnings in cone contrast processing.

To help bring this out, first notice that just as differences in cone(-opponent) contrast predict differences in colour appearance (at least under simplified viewing conditions), equivalent cone contrasts should predict sameness in colour appearance, assuming other factors are held constant. The latter hypothesis is known as the *cone contrast rule* (Whittle, 2003). Figure 2 illustrates a case in which this is borne out quite strikingly, if only approximately. In 2(a), the squares in the top and bottom rows were carefully chosen to

¹⁴ For example, see Conway (2001).

¹⁵ Byrne & Hilbert (2003) argue from colour constancy to reflectance physicalism about colour. Allen (2016: Ch2) argues from constancy to the mind-independence of colour, addressing various critical responses to Byrne & Hilbert.

produce the same relative cone responses with their respective backgrounds. That is, in column 1, the cone contrasts between the top square and the grey background are identical to the contrasts between the bottom square and the green background. Now, although there are some discernible differences in the appearance of these squares, they appear very similar. Compare this with 2(b), which shows the exact same squares as 2(a) – that is, corresponding squares have the same absolute colour values – but all against a uniform grey background. The two rows now appear quite different. For me, this is especially marked in columns 3 and 4.




Figure 2, reprinted with permission from Hurlbert (1996)

Let's reflect on the phenomenology of image 2(a) a bit more. Although the display is artificial and contains no explicit cues as to the illuminant, a natural interpretation is that 2(a) displays the same row of squares under two different lights: a fairly neutral light on the top row, and a greenish light on the bottom row. Alternatively, on the bottom row, it might seem as if one is viewing the same set of squares through a green transparency. This is no accident. As it turns out, for almost every change in illumination that we ordinarily encounter, the corresponding change in cone excitations preserves spatial contrasts within each cone class.¹⁶

¹⁶ Introducing a coloured transparency also tends to preserve such contrasts. See Westland & Ripamonti (2000).

That is, for natural illuminants A and B , $L(p)/L(q)$ under $A \approx L(p)/L(q)$ under B . This is the *illumination-invariance of cone contrasts*, or *invariance of cone excitation ratios* (Foster & Nascimento, 1994; Foster, 2011: §5.3). Significantly, cone-opponent contrasts exhibit the same, if not slightly higher degrees of illumination-invariance.¹⁷ That is, for example, $|RG(p)-RG(q)|$ under $A \approx |RG(p)-RG(q)|$ under B . This is one of the most profound features of the visual processing of colour. To illustrate, each point in Figure 3 represents a pair of excitation ratios within a single cone class ($a = L$, $b = M$, $c = S$), between two randomly chosen surfaces in a natural scene, viewed under skylight and sunlight. These ratios fall very close to the diagonal, which represents perfect equivalence across the illuminant change. The pattern holds for any two illuminants drawn at random from the daylight set, and with objects drawn from a sample of 640,000 reflectance spectra gathered from natural scenes (Nascimento et al., 2002).

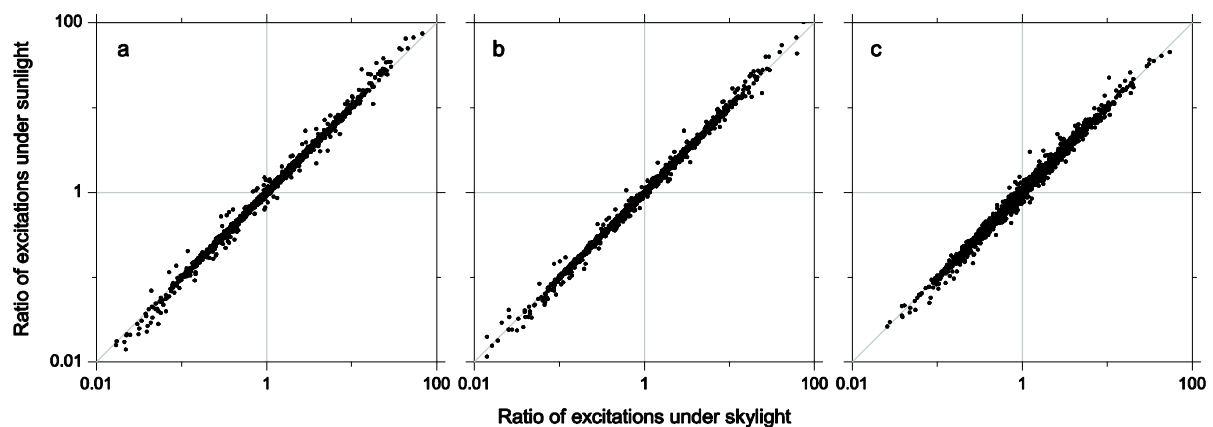


Figure 3, reprinted with permission from Foster (2011)

The almost perfect illumination-invariance of cone contrasts has a corollary in human behaviour. In a paradigm developed by David Foster and his collaborators, subjects are

¹⁷ Foster (2011: 692), Foster et al. (2000: 181), Linnell & Foster (1996: 226), Zaidi (1998: 1772), and Zaidi et al. (1997).

presented with a reference stimulus, alongside (simultaneous) or followed by (successive) a test stimulus that either preserves ('illumination condition') or violates ('surface condition') the cone excitation ratios in the reference stimulus.¹⁸ Subjects are asked to judge whether the difference/change they encounter is a difference/change in illumination, or in the surface material properties of the stimulus. We are exquisitely attuned to this difference, routinely judging a change that preserves cone ratios to be a change in illumination, and a change that violates these ratios as a change in surface properties. This discriminatory capacity is usually taken to operationalise relational colour constancy, discussed above. To reiterate, then, it is no accident that we see image 2(a) as presenting the same set of squares under two different illuminants, given that the cone contrasts between squares and backgrounds are the same in the two rows.

On my framework, then, some types of colour constancy are conceived as corollaries or consequences of the illumination-invariance of cone contrasts. To be clear, I do not think that this framework is apt to explain all types of colour constancy. One lesson of pluralism is that we should not expect a one-size-fits-all account. Nonetheless, the framework seems apt to explain so-called 'appearance-based' or 'phenomenal' colour constancy,¹⁹ at least in part, and relational colour constancy. This framework has a number of advantages, in the present context.

Firstly, it highlights a basic computational means of distinguishing changes in proximal stimulation that are due to the illumination, from changes that are due to surface colour properties. As noted in section 2, regarding relational colour constancy, the associated

¹⁸ The seminal works are Craven & Foster (1992) on successive constancy, and Foster et al. (1992) on simultaneous constancy.

¹⁹ See Davies (2016) and Wright (2013).

mechanisms need only compute spatial or temporal ratios of cone excitation, and assess whether these ratios are invariant, (in which case the variation/change is likely due to the illumination), or variant (likely due to a variation/change in surface colour). Mechanisms such as these seem promising, from the perspective of limning the most basic kinds of colour constancy.

Secondly, these mechanisms subserve a robust, well-documented perceptual capacity: the ability to discriminate, on the basis of vision, between differences/changes in illumination, and differences/changes in surface colour properties. I noted above that this discriminatory capacity is usually taken to operationalise relational colour constancy. Which is to say, this capacity is not necessarily grounded in any invariance in one's experience of the 'absolute' colour values of objects. In fact, as also noted above, the actual perceptual bases for these discriminations are not yet clear: I tendered only that such discriminations indicate that one *in some sense* perceives the colour relations between objects as invariant under illumination changes, and variant under surface colour changes. This phenomenological neutrality is advantageous here, as it leaves open the question of how information about colour relations may be conveyed to one in visual experience. Such information may, of course, be conveyed in different ways, at different times. The framework thus allows for productive speculation as to whether, and if so, how such information may be conveyed even to an achromatopsic subject, despite their inability to perceive the colours of things. To pre-empt what lies ahead, my view is that awareness of qualities of the edges of surfaces may serve just this function.

I now set this framework into action, in addressing our puzzle. Crucially, despite being severely impaired in their conscious experience of colour, some complete achromatopsics are able to discriminate illumination changes from surface colour changes, indicating visual sensitivity to preservations and violations, respectively, of cone excitation

ratios. I shall now set out the details of one such case, and argue that it constitutes a genuine case of colour constancy, albeit of an extremely unusual and impoverished kind.

Hurlbert and colleagues (1998) tested MS using a simultaneous asymmetric colour-matching task, involving two types of stimuli. The *simple* stimuli were uniformly coloured square patches set against similarly uniform coloured backgrounds. The *complex* stimuli involved the same coloured patches set against coloured Mondrian backgrounds. The reference patches were greenish-yellow, blue, blue-green, and purple, with additional yellowish-green and orange patches used in some conditions. For simple stimuli, the reference background was a neutral grey under a simulated daylight illuminant. Test patches were generated by shifting the cone coordinates of the reference patches in one of three ways: an $L+M$ shift, $L-M$ shift, or S shift. The same shift was applied to the reference background, so that test patch-background pairs had the same cone ratios as the corresponding reference stimulus. MS was tested using a single interval forced-choice paradigm. Each test stimulus was presented alongside one of two possible matching alternatives: an ‘incorrect’ match, or a ‘correct’ match. The ‘correct’ match would be the original reference patch-background pair from which the test was derived, which therefore had the same cone contrasts. The ‘incorrect’ match would be the test-shifted patch against the original reference background. This alternative patch therefore had the same absolute cone coordinates as the test patch, but different cone contrasts. The task was to say whether the two stimuli were the same or different. As intimated by Hurlbert and colleagues (1998: 136), the task effectively probes whether MS’s responses fit the normal pattern predicted by the cone contrast rule, which manifests a ‘form of colour constancy’ (1998: 137), for reasons discussed above.

Surprisingly, MS displayed near normal discrimination performance for both achromatic $L+M$ and chromatic $L-M$ shifted stimuli. His d -prime for the $L-M$ stimuli was 0.98, or 69% correct. For $L+M$ stimuli, his d -prime was around 2.4, but this difference was

largely attributable to the fact that the average cone contrast distance for $L+M$ stimuli was around four times larger than for $L-M$ stimuli. MS was unable, however, to discriminate differences in S -cone contrast alone. Significantly, MS was also completely unable to discriminate chromatic contrast with complex stimuli. I return to this point later.

One immediate worry is that MS is perhaps picking up on differences in luminance, or luminance contrast between the correct and incorrect matching alternatives. Perhaps contrasts of this sort produced discriminable differences in the perceived lightness – that is, the achromatic colours – of the stimuli, of which MS was aware. But there were no differences in luminance, hence luminance contrast, between the correct and incorrect matching alternatives: they differed solely in chromatic or $L-M$ contrast (1998: 138). A further control experiment confirmed that MS could not have been responding to luminance contrast (1998: 141). Other studies confirm that MS is able to detect the presence of features defined solely by chromatic contrast, despite random temporal or spatial luminance masking (Barbur et al. 1994; Heywood et al., 1994).

A second worry is that the ‘correct’ matches in this experiment might not be those chosen by a healthy, colour constant subject. In a further control, however, a healthy subject completed an asymmetric colour matching task with the same test patch-background combinations, using matching-by-adjustment rather than forced-choice discrimination. The healthy subject’s preferred matches fell ‘almost exactly on the equal-cone-contrast line,’ (1998: 142), although there was more deviation from perfect S -cone contrast. That is, the points at which test and reference patches appeared the same in respect of hue, saturation, and luminance, closely corresponded to the points at which the patches had equivalent cone contrasts with their respective backgrounds. Accordingly, in successfully discriminating contrast-preserving from contrast-violating alternatives, MS is discriminating stimuli that exhibit high versus low degrees of colour constancy for a healthy subject.

These points, while important, largely reflect concerns about the empirical details. A more philosophically minded worry is that, even taken at face value, the data only establish a behavioural or extensional equivalence between MS and healthy, reasonably colour constant subjects. Such equivalence does not suffice to show that the perceptual basis for MS's discriminations constitutes a genuine type of constancy. After all, when healthy subjects perform well on such tasks, they do so by attending to the colours of the stimuli. Their colour constancy does not lie in the discriminatory behaviour: it lies in the perceptual response, in how they see things colour-wise. Interesting as it may be, the objection runs, MS's discriminatory behaviour is not really the issue. Either MS is completely colour blind, as advertised, in which case he could not possibly have colour constancy in the usual sense; or MS is not completely colour blind after all, in which case he is presumably just discriminating these stimuli on the basis of how they look to him colour-wise.

Let me start with the second horn of this objection. It has been suggested that MS might have some residual visual experience of colour, but of a degraded or unusual sort. For example, MS might see 'partial colours,' qualities defined by at most two of the standard three dimensions of hue, saturation, and lightness.²⁰ Alternatively, MS might see 'alien colours,' properties located in quality spaces quite different to our own hue, saturation, and lightness space.²¹ It is impossible to rule out these possibilities completely. In other work, however, I have argued that neither proposal is especially plausible (Davies, 2021). In any event, in the next section I present what I take to be a more probable explanation.

I shall therefore focus here on the first horn of the objection. My first point is that colour constancy in the 'usual sense' presumably entails a capacity to perceive colours, in

²⁰ Brown (2014: 14ff).

²¹ MacPherson (2015).

line with Constancy for Colours. In arguing for pluralism, however, I have already established that this concept does not exhaust the possible forms that colour constancy may take. The real issue, as I see it, is whether MS has colour constancy in *any* sense. And admittedly, thus far, the only real evidence for this claim is that MS performs reasonably well on a standard psychophysical colour constancy task. What further reasons can be supplied?

My first argument focuses on the relationship between MS's discriminatory abilities and the environmental conditions involved in normal cases of colour constancy. In summarising the constancy experiments, I claimed that MS successfully discriminated stimuli that either preserved or violated achromatic ($L+M$) or chromatic ($L-M$) cone excitation ratios. Considered in the void, such discrimination bears no obvious connection to colour constancy. Constancy involves the discrimination of distal conditions, involving surface colour and illumination, not proximal conditions, such as cone excitation ratios.²² Recall, however, that changes in daylight illumination almost flawlessly preserve such ratios, whereas changes in surface colour properties routinely change them (Foster & Nascimento, 1994; Foster, 2011: §5.3). As such, the capacity to discriminate stimuli that preserve or violate such ratios is, ipso facto, a capacity to perceptually discriminate illumination changes from surface colour changes. But if a subject is capable of perceptually discriminating illumination changes from surface colour changes, then her visual system must be processing the incoming perceptual signal in a way that distinguishes changes in the signal that are likely due to changes in the illumination, from those likely due to changes in surface colour. Any visual system that can do these things, in my view, is a system that exhibits colour constancy.²³ For it is plausibly a minimal, core part of any understanding of colour constancy, that it requires possessing a

²² Cf. Burge (2010: 396ff) for related points.

²³ I provide a lengthier defence of this claim in Davies (2018).

visual system that can distinguish or disentangle the contribution of the illumination and surface colour to the incoming perceptual signal, and produce therefrom some sort of differential perceptual response to the two types of environmental condition. MS's visual system does these things; hence, his discriminatory ability suffices for colour constancy.

My second argument is that the neurological evidence suggests that MS retains some key mechanisms of colour constancy – though he clearly lacks others. These include retinal mechanisms and, crucially for my view, early cortical mechanisms. Regarding the former, given his intact retinal cone systems (Mollon et al., 1980), MS likely possesses retinal mechanisms of light adaptation, such as von Kries scaling. Retinal light adaptation is universally thought to make a significant contribution to appearance-based or phenomenal colour constancy, involving invariance in respect of hue-saturation content. When experimental conditions are conducive to high levels of adaptation, phenomenal constancy is correspondingly robust.²⁴ In this connection, it is significant that MS underwent dichoptic presentation of the test and reference stimuli, with an adaptation period of two minutes to the respective background illuminants. This would have engaged retinal adaptation mechanisms, if MS had them. These mechanisms therefore plausibly form at least part of the explanation of MS's discrimination performance, although of course it remains unclear, for now, how such activity might shape his visual experience, in the absence of hue-saturation content.

It is unlikely, however, that retinal adaptation mechanisms were the sole driver, perhaps even the main driver of MS's performance. Recall that MS was only able to discriminate cone contrasts for simple stimuli. When stimuli were presented against complex Mondrian backgrounds, MS's performance effectively dropped to zero for chromatic *L-M* shifts (1998: 142). Now, MS was pre-adapted to the simple, spatially uniform, reference and

²⁴ Kuriki and Uchikawa (1996).

test backgrounds for the same period in both the simple and complex stimulus conditions (1998: 138). As such, cone responses to the reference and test patches should have been rescaled to approximately the same extent in each condition (Hurlbert & Wolf, 2004: 157; Kentridge et al., 2004: 828). MS's struggles with complex stimuli therefore imply that some other factor, in addition to adaptation, contributed to his discrimination of simple stimuli.²⁵

It is therefore likely that MS retains some cortical mechanisms of constancy, in addition to retinal mechanisms. These putative cortical mechanisms are likely involved in computing cone(-opponent) contrasts. Areas V1 and V2 are preserved in MS's left hemisphere, as evidenced by the conscious vision in his right visual hemi-field. V1 is destroyed in the right hemisphere, accounting for his blind left visual hemi-field. The damage to lingual and fusiform gyri, which characterises cerebral achromatopsia, is bilateral. MS therefore retains the machinery of double-opponent cells in V1, for example, which are widely taken to be involved in coding chromatic contrast across spatial boundaries (Gegenfurtner & Kiper, 2003; Hurlbert, 2003; Kentridge et al., 2004: 821, 829). Given the aforementioned connection between contrast and constancy, Hurlbert and Wolf (2004: 147) argue that 'V1 and lower areas may therefore play a larger role in colour constancy than previously thought.' Conway & Livingstone (2006: 10842) suggest similarly that double-opponent cells in macaque primary visual cortex provide 'ideal substrates for colour constancy and colour contrast..., computations that likely involve V1.' Using fMRI, Barbur and Spang (2008) found that rapid or 'instantaneous' colour constancy in normal subjects produced strong activation in areas V1, V2, and V3, in addition to the fusiform colour area

²⁵ Kentridge et al. (2004: 828) also report that MS could discriminate invariant from variant local cone contrasts 'in free viewing with a series of relatively rapidly changing displays,' which 'clearly depends upon processes beyond simple adaptation.'

V4. Importantly, using the same paradigm, Barbur and colleagues (2004) found evidence for reasonable degrees of instantaneous constancy in three achromatopsic patients, whose lesions at least partially covered V4. They conclude that ‘the evidence... points to V1, ... as the principal neural substrate for mediation of instantaneous colour constancy,’ (2004: 27).

Let me summarise the argument so far. My philosophical framework ties some types of colour constancy to the illumination-invariance of cone contrast coding. Deploying this framework, I argued that MS displays a key hallmark of colour constancy, by virtue of discriminating between simple stimuli that maintained or violated cone contrast ratios, which amounts to discriminating distal changes either in illumination conditions, or surface colour. I further argued that MS retains some mechanisms of colour constancy, including cortical mechanisms implicated in the computation of cone(-opponent) ratios. In my view, these points provide cogent grounds to believe that MS possesses a kind of colour constancy. I close this section by highlighting some wider implications of this view.

One distinctive feature of my view is that it pinpoints the close connection between colour constancy and contrast. As noted above, philosophers have tended to regard these phenomena quite differently. One notable exception, however, is Akins and Hahn’s (2014) timely paper, which argues that the colour vision system is best understood as comprising different contrast systems, working in concert to facilitate object seeing. Akins and Hahn target the traditional conception that colour vision is primarily for *seeing colours*; that is, for perceiving the sensuous colour qualities of objects. Relating this to my work, the traditional conception clearly undergirds the conception of colour constancy provided by Constancy for Colours. On this conception, colour constancy with respect to X (for S , at t) is defined as constancy in respect of the colour property that S ’s visual experience at t presents X as having. That is, satisfying Constancy for Colours amounts to having *constancy for seen colours*. As Akins and Hahn point out, however, the traditional conception of colour

perception does not mesh well with the contrast-driven nature of chromatic processing. They ask, what is contrast processing really good for? The answer, they contend, is that contrast processing is very good at helping us see *things*. In particular, it is very good for marking out and emphasising changes in the perceptual signal that are likely to correspond with the borders of objects, and distinguishing these from changes that correspond with mere fluctuations in illumination.

Now, although Akins and Hahn do not discuss constancy explicitly, the connections between their view and mine should now be coming into view. Firstly, while Akins and Hahn emphasise the importance of contrast processing for colour vision quite broadly, my framework brings out its fundamental importance for colour constancy. Secondly, we both want to move the philosophical discussion away from traditional, intuitive conceptions of colour perception (constancy), which focus on our conscious awareness of sensuous colour properties, or ‘seeing colours’.²⁶ Thirdly, and finally, Akins and Hahn suggest that colour vision is primarily for seeing *things*; that is, for perceptually differentiating or singling out ordinary objects. As we shall see imminently, my account leads us in a similar direction. Specifically, I shall argue that the putative contrast-based constancy exhibited by MS involves an invariance in the visual perception of the *chromatic edges* of objects, across changes in variations in illumination. As such, building on Akin and Hahn’s stimulating work, my view extends and precisifies the sense in which colour perception enables us to see things.

²⁶ See also Davies (2018).

4. The Chromatic Edge Hypothesis

Thus far, I have argued that an achromatopsic patient possesses a genuine type of colour constancy, despite the profound impairment of their conscious experience of colour. This account, however, contains a significant lacuna: absenting any awareness of colour, what *is* the conscious upshot of this putative constancy capacity? This is the most puzzling aspect of the paradox of colour constancy. I now present the chromatic edge hypothesis, which holds that the constancy in question involves the conscious perceptual representation of chromatic contrast properties ‘at’, in a sense to be explained, the edges of surfaces.

To begin, recall from section 1 that achromatopsics have three fully functional retinal cone mechanisms, and retain cone-opponent processes. Although they are colour blind by normal standards, achromatopsics are able to recruit this preserved colour pathway in some surprising ways. To pick one example, MS could detect a coloured square concealed in a grey checkerboard, so long as the colour was maximally saturated and the background luminance contrasts were small (Heywood et al, 1994: 252). As I have already noted, these abilities seemingly involve conscious perception: achromatopsia is not blindsight for colour.

Heywood and colleagues (1998: 413) report that ‘in no case, is a patient required to “guess” the identity of an invisible figure concealed in’ such displays. Rather, ‘residual processing, of whatever origin, resulted in a conscious perceptual change, notwithstanding the absence of colour qualia,’ (1998: 415).

How is this possible? The most plausible interpretation is that achromatopsia involves a dissociation between two perceptual functions normally fulfilled by chromatic processing in the visual system: so-called ‘colour for colouring’ and ‘colour for form.’ Healthy subjects have almost identical chromatic contrast thresholds for detecting a difference in colour between two abutting surfaces, and detecting the presence of a contour or edge between them.

That is, the magnitude of chromatic contrast required for the subject to see a difference in colour between two regions is the same as that required to see an edge between them. For this reason, it is hard for us to conceive of a visual experience of contour that was not accompanied by an appreciable difference in colour. Indeed, it is generally considered necessary a priori that our visual experience of differences in colour *determines* our visual experience of such demarcations of form. This idea goes back at least to Aristotle (1991: 3),

The faculty of seeing, thanks to the fact that all bodies are coloured, brings tidings of multitudes of distinctive qualities of all sorts; whence it is through this sense especially that we perceive the common sensibles, viz. *figure, magnitude, motion, number...*

As Sorabji (1971: 61-2, fn.27) explains,

If there is a sharp boundary between an area of one colour (hue, saturation, or brilliance) and an area of another, and if we see where the boundary runs, this *is* to see (part of) the shape of the areas.

Achromatopsics, however, display marked differences in these contrast thresholds.²⁷ As one would expect, their thresholds for detecting colour differences are significantly raised. Also in line with predictions, patients show higher thresholds for detecting chromatic contour. Surprisingly, however, colour thresholds are raised significantly more than contour thresholds. In other words, although achromatopsics are worse than healthy subjects at detecting chromatic contours, they are much better than one would predict, given the extent of their impairment in detecting differences in colour. Achromatopsics seem capable of using

²⁷ See Barbur et al. (1994), Mollon et al. (1980) and Heywood et al. (1991).

information from colour contrasts to perceive form, despite being unable to use it to perceive different colours.

Detailing this view, Barbur and colleagues (1994: 332) suggest that chromatic signals can have at least two distinct functions, and... these functions can be affected differentially by the lesion [in achromatopsia]. [These] signals carry sufficient information to enable the generation and spatial representation of an object in terms of its form and structure, and they can be used to generate at least one more visual attribute, namely the perceived object colour.

Kentridge and colleagues (2004: 822) concur:

the distinction between wavelength processing for the assignment of surface colour and for segmenting chromatic boundaries is not apparent in the normal observer... However, these wavelength-based processes are neatly dissociated in cases of cerebral achromatopsia.

Similarly, Chirimuuta and Kingdom (2015: 226):

it seems reasonable to conclude that these cerebral achromatopsics have a selective loss in the ability to use colour vision to see colours (experience hues), but not a loss in the ability to use colour vision to see form.

Clearly, the colour for form interpretation is quite widely endorsed.

The interpretation provides the best explanation of the surprising colour-related abilities of achromatopsic patients. One early finding was that some patients could detect figures in Ishihara colour plates at a distance of two metres, though not at normal reading distance. As Mollon and colleagues (1980: 133) observed, at two metres 'the luminance contours of individual discs are no longer resolved and the dominant contour is the hue

boundary between figure and ground.’ It seems that the achromatopsics consciously perceive this chromatic contour – they can name the figure and trace its outline – despite the discs appearing random shades of grey. In another study, MS was able to discriminate sequences of isoluminant squares ordered in respect of colour, from sequences of randomly ordered squares. Heywood and colleagues (1991: 802) note that his ‘verbal replies showed that he did so by detecting an edge between two stimuli that were, to him, perceptually identical.’ Heywood and colleagues (1994: 252) found that MS could also detect a highly saturated coloured square in a grey checkerboard of varying luminance, seemingly because ‘saturated chromatic and achromatic boundaries are conspicuously different to M.S., particularly in dynamic displays, when they are of similar luminance contrast.’

Importantly, MS is not just able to detect the presence of chromatic edges: he can discriminate differences in type of chromatic edge as well; types that are individuated by differences in the magnitude and direction of chromatic contrast that occurs at the boundary. For example, MS could reliably discriminate the odd one out in displays such as Figure 5(b) and 5(d). These images show three coloured discs against an equiluminant background, where the odd one out is defined either by a different magnitude of cone-contrast from the background, as in (b), or an equal-magnitude contrast in a different direction in colour space, as in (d). Again, it seems that MS does so by attending to differences in the visual appearance of the edges of the discs, which otherwise appear to him to be identical.²⁸

²⁸ Kentridge et al. (2004: 828) suggest that this is explained by the activity of double-opponent cells, which are ‘good candidates for the mediation of the perception of form from colour.’ Tellingly, they add that these cells are ‘a likely contender for an early stage in the maintenance of colour constancy,’ and ‘are well-suited to segmenting the visual scene on the basis of colour variation,’ (2004: 822).




Figure 5 (Reprinted with permission from Kentridge et al. 2004)

The chromatic edge model aims to explain the nature of such unusual visual experiences. I assume that when we visually perceive an object, we represent its surfaces and their edges. I further assume that such representations have an iconic format. Now, in ordinary cases of object seeing, where an object is seen as a figure against its background, the edge boundaries of its surfaces appear to be owned by the figure, rather than the ground. In the topological jargon, figures appear ‘closed,’ while grounds appear ‘open,’ in the region of their boundary with the figure. This view was clearly articulated by the Gestalt psychologists. Koffka (1936: 181), for example, reports that in perceiving a figural object, he perceived a ‘contour line’ which ‘*belonged* to the enclosed figure and segregated it from the surrounding field.’ Elsewhere, Goldman’s (1977: 280) important ‘differentiation condition’ on object seeing requires the ‘representation of an edge or boundary.’²⁹ Casati and Varzi (1999: 71) claim similarly that boundaries are ‘*bona fide* spatial entities’ that ‘enter the content of our perceptions.’

Now, iconic representations are *holistic*, in that they lack canonical decompositions into syntactic constituents that separately denote individuals and features. Take a picture of Boris Johnson. There are no constituent parts of this image that separately represent parts of Boris Johnson and the features of those parts, such as their shape, size, and texture. Rather,

²⁹ Cf. Burge (2014: 492, fn.7).

any constituent part of the image will jointly represent a part of Boris, together with the features of that part. Let's assume that this image has two types of primitive: there are *surface region primitives*, which represent parts of the surface of Boris, and *edge primitives*, which represent segments of the edges of Boris's surfaces. Surface region primitives holistically encode size, shape, texture, and monadic colour information. Edge primitives, in contrast, holistically encode curvature, length, orientation, and – crucially, for my view – chromatic contrast information. We can think of these edge-based chromatic contrast contents as directed magnitudes in colour space: vectors from one location (though no particular location) in colour space to another. The vector might tell us, for example, that an edge marks a contrast of $+N$ units along the red-green axis, or $-M$ units along the yellow-blue axis. This leaves open the spatial orientation or 'polarity' of the contrast.³⁰ In Figure 5(d), for example, the boundary between the top red disc and the yellow background marks a strong +red contrast in the direction of the disc. To accommodate this, recall that on the Gestalt view of figure-ground, edges are represented as belonging to just one of the surfaces that they bound. In this example, the circular boundary is represented as belonging to the disc, which appears 'closed,' rather than its background, which appears 'open'. This edge-ownership feature determines the spatial polarity of the +red contrast: the edge representation 'says' that things get redder by a certain magnitude, in the direction of the surface that owns the boundary.

My view is that this chromatic edge content partly determines the phenomenal character of our visual experience of surface boundaries. Edges look to us in certain ways in visual experience, and these appearances vary according to the chromatic composition of the scene, among other things. These differences in appearance supervene on differences in

³⁰ On the neural representation of chromatic contrast polarity and border ownership, see Friedman et al. (2003). See also Davies (2020: §3.3).

chromatic edge content, that is, in the spatially oriented, directed magnitude of colour contrast represented at that location. This predicts that our experience of an edge between a red and yellow region will have a different character to our experience of an edge between a red and green region, or a red and blue region, assuming that we control for luminance contrast. It is an open question whether chromatic edge appearance can be modelled in something like the ordinary opponent-colour space, but I assume not. More likely, the quality space for edges will be of lower dimensionality, *sui generis*, and will depend in complex, context-dependent ways on both luminance and chromatic contrast.³¹

The chromatic edge model is independently motivated by a range of evidence (Davies, 2021). The present aim is to apply the model to theorise the unusual constancy capacity in achromatopsia. The standard interpretation of achromatopsia as involving preserved colour-for-form function provides a presumptive reason to favour my account. In what follows, I argue more directly in its favour. To begin, recall that changes in illumination almost flawlessly preserve spatial (and temporal) cone (and cone-opponent) contrasts. The flipside is that changes in surface colour routinely change these contrasts. As noted above, it is plausible that our almost unerring capacity to discriminate illumination changes from surface colour changes is a corollary of this pattern. Now, the illumination-invariance of cone(-opponent) contrasts likely serves a dual purpose in early vision. Via the colour for colouring function, these invariant contrasts provide a partial basis for computing (approximately) invariant representations of surface colour, in line with the cone contrast

³¹ As Akins (2014: 200-201) notes, luminance edges and chromatic edges are largely statistically independent in visual images. As such, having a distinct capacity to represent chromatic edges *per se* would greatly increase our chances of discriminating objects from their backgrounds.

rule. In addition, these contrasts plausibly input into edge computations, via the colour for form function. Suppose that a red disc viewed against a yellow background under direct sunlight at $t1$ produces a +red opponent contrast of N . If the light changes to skylight at $t2$, the contrast between disc and background will remain at roughly N +red. By hypothesis, these contrasts drive the formation of representations of the disc's boundary with its background. The invariance in the chromatic contrasts registered at $t1$ and $t2$ determine, along with registrations of the boundary's orientation and length, that these representations categorise the edge as being of an invariant spatial and chromatic type. In particular, the edge content will attribute a constant chromatic contrast of N +red across this temporal interval. This invariance in chromatic edge content, I suggest, determines a constancy in the visual appearance of the boundary across the change in illumination. The edge of the disc, hence the very form of the disc, has an invariant appearance under changes in illumination. It consistently looks the way that the boundary of a red object against a yellow background looks, which I contend is different, for example, from the way the boundary of a blue object against a yellow background looks.

To be clear on the implications here, the claim is that the colour for form function, just as much as colour for colouring, has associated perceptual constancies. Recalling Akins and Hahn (2014), the former constancies are associated with our capacity to see things, while the latter, more familiar constancies are associated with our capacity to see colours. The former constancies involve an invariance with respect to the chromatic contrast properties that are attributed in perceptual representations of surface boundaries. This representational invariant constitutes a primitive type of colour constancy, a basic objectifying capacity within the colour vision system.

Although this capacity undoubtedly falls outside our ordinary conception of colour perception, it is nonetheless intuitive that we have it. The familiar fact is that when we look at

an object under changing light, we find that its absolute colour tends to appear fairly constant. Intuitively, however, we also find that our awareness of the form of the surface – the sense in which the surface appears as demarcated or differentiated from its background – exhibits its own, distinctive sort of constancy. The aspect of experience that I am drawing attention to is not our awareness of the surface's shape per se; of its circularity, say. Rather, it is our awareness of the limits of the surface, its outermost bounds; the elusive, widthless parts that determine its shape properties, as opposed to constituting them. It seems intuitive, for example, that a change in lighting typically would not result in a surface appearing much more (or less) pronounced, in terms of how strongly (or weakly) our visual experience differentiates or segments the surface from its background. Indeed, it seems the Gestalt psychologists knew this (West et al., 1996), but as often happens, things that should not have been forgotten were lost.

We can speculate on why such chromatic edge constancies should exist. It has been suggested that colour perception might have evolved to facilitate the detection and segregation of coloured fruits against dappled foliage (Mollon, 1989). The opponent axes that characterise human colour processing are well suited to emphasising strong red-green contrasts, as with a red berry in a leafy bush, which are conveniently orthogonal to contrasts along the yellow-blue axis, about which most natural changes in light occur. Now, one way to achieve good segregation would be to implement a colour for colouring function that produces representations of absolute surface colour properties, which stand in the relevant contrast relations. Constancies associated with this function would ensure that the strong contrast between the absolute redness of the berry, and the absolute greenness of the leaves, is not diminished by changes in illumination. This would be one way – and a very good way – to ensure that the berry segregates well from surrounding foliage, even in dappled conditions. But it need not be the only way, nor even the most basic way to achieve this aim.

Here is an alternative formula: first, use strong chromatic contrasts to generate representations of surface edges; second, produce different types of edge representation, depending on the particular direction and magnitude of chromatic contrast; third, ensure that these edge representations exhibit high levels of invariance under changes in illumination. By this formula, the strong opponent red-green contrast between the berry and leaves would produce an edge representation of a different type to those produced by the weaker, non-opponent yellow-green contrasts between different leaves. The edge of the berry would therefore appear different to all the other edges in the vicinity, despite the many criss-crossing luminance variations caused by the dappled illumination. Intuitively, the berry would segment more strongly, its form appearing more pronounced, relative to its background, than any of the leaves, relative to their surrounds. Crucially, moreover, the berry would appear no less pronounced were a cloud to pass overhead, the dominant illumination changing from sunlight to skylight, given the illumination-invariance of the contrasts that drive the edge representations. Such a capacity might well facilitate the detection and segregation of coloured fruits in dappled conditions. Indeed, I can imagine creatures unable to perceive the absolute colours of surfaces, but nonetheless able to forage successfully using only their awareness of such boundary-based contrasts, alongside capacities for size, shape, and distance perception.³² These contrasts would bring red berries to the fore in their visual

³² Compare Mollon's (1989: 21-26) discussion of the difficulties in foraging encountered by subjects with acquired achromatopsia. Although these subjects can no longer distinguish berries 'by the Colour from the neighbouring Grass,' they can distinguish them 'by the shape' at close distances (1989: 23). We can speculate, however, as to whether such detection is at all facilitated by the representation of strong chromatic contrasts at the boundaries of the berries, as proposed here, or whether it is purely a matter of shape-based detection.

experience of the scene, despite their experience lacking any sensuous chromatic phenomenology. For such creatures, chromatic edge constancy would ensure that the berry segregates consistently across all phases of daylight, and other, more transient changes in illumination.

Moving on, an important advantage of the chromatic edge hypothesis is that it would explain the drop-off in MS's performance between simple and complex stimulus conditions, where retinal adaptation could not. Recall that MS could discriminate cone contrasts for patches against uniformly coloured backgrounds. In these conditions, there were uniform chromatic contrasts along the entire boundary of each patch. In the complex condition, patches were embedded in Mondrian backgrounds composed of rectangles of different colours. Each patch therefore abutted several differently coloured rectangles in its background, producing non-uniform chromatic contrasts along its boundary. Now, the chromatic edge model predicts that in the simple condition, the uniform chromatic contrast between a test patch and its background should result in its edges having a uniform visual appearance. In the complex condition, the varying chromatic contrasts along the perimeter of the test patch would produce variations in the visual appearance of this boundary. This predicts that it would be easier to discriminate between 'correct' and 'incorrect' matching alternatives in the simple condition, than in the complex condition. On the chromatic edge model, a 'correct' match would have edges with the same overall visual appearance as the edges of the test patch. An 'incorrect' match would have edges that appeared overall different to the edges of the test patch. Presumably, it would be easier to assess these options between patches with uniform contrasts, hence uniform appearance along their entire boundaries. In the simple condition, if the edge of the test patch looks uniformly F , and the edge of the matching alternative looks uniformly G , then the subject may easily infer that they are different. In contrast, in the complex condition, if the edge of the test patch looks in part F , in

part *G*, and in part *H*, and the edge of the matching alternative looks in part *F'*, in part *G'*, and in part *H'*, then it would take a more complicated comparison to discern whether they are different (as indeed they are).³³

In noting this advantage, however, we have uncovered a potentially worrying implication of my view. A change in background colour alone will produce a change in the chromatic contrast between the patch and its background. As such, if we take a patch (or patches) of constant colour and present it successively (or simultaneously) against differently coloured backgrounds, my view predicts that the chromatic edge content of our perceptual representation of the patch will change accordingly. Chromatic edge content therefore will be highly variant under changes/variations in background colour. This implies, however, that chromatic edge content cannot help explain so-called *background-independent colour constancy*, however well-suited it seems to explaining *illumination-independent colour constancy*.³⁴ This is potentially problematic, because in normal viewing conditions, we perceive objects as fairly constant in colour across both types of change. It seems reasonable therefore to expect our ultimate theory of colour constancy to explain both types of case.

This objection raises an important and neglected question regarding the scope of theories of colour constancy. Almost without exception, philosophical discussions of colour constancy have focused only on the illumination-independent kind.³⁵ Much the same is true of the science. Failures of background-independent colour constancy are frequently

³³ I use *F*, *G*, ... as placeholders for the different looks that edges might have, as a result of their different chromatic contrasts.

³⁴ This distinction originates in Whittle and Challands (1969: 1108).

³⁵ Allen (2016: Ch2) considers both types, and seems to want to account for both within a single theory of colour constancy.

discussed, but not under that description: they are usually classified as colour contrast effects, as discussed in section 3. That we in fact have robust background-independent constancy is typically overlooked. Now, I fully acknowledge the extent of our background-independent constancy. All the same, my view only explains one type of illumination-independent constancy. The issue is whether this partitioning of illumination- and background-independent constancy is appropriate.

My view, building on the pluralism of section 2, is that each type of constancy needs independent theoretical treatment. At the end of inquiry, we may expect that certain commonalities or overlapping features will emerge. As a research programme, for now, however, I believe that we should embrace splitting rather than lumping. The splitting approach is supported by considering the computational descriptions of these constancies. As I assume will be familiar, on Marr's influential view, the computational level description of a cognitive capacity should specify the 'why' of the computation, among other things. At a suitably fine-grained level of detail, the 'why' of illumination-independent constancy clearly differs from the 'why' of background-independent constancy. The former enables organisms to derive stable representations of colour-related properties of objects, despite occupying environments in which the illumination conditions are continuously varying across space and time. The latter, in contrast, enables organisms to derive stable representations of colour-related properties of objects, given that they are (presumably) mobile creatures with varying spatial perspectives on the world, and given that many objects of perception are themselves mobile, hence liable to be viewed across changes in background conditions.

Now, a lumper might suggest that these computational descriptions could be conjoined, yielding a combined illumination-and-background-independent conception of

colour constancy.³⁶ But remember what the computational level description is supposed to do within the theory. For one, it is supposed to inform and constrain the search for appropriate algorithmic- and implementation-level descriptions of the capacity. On this score, the splitting approach has proven very productive. For it seems that the visual system does contain information-processing mechanisms specifically geared towards resolving the illumination-independent constancy problem, and which are ill-suited to resolving the background-independent constancy problem. As Richard Brown (2003: 248) notes,

There is an interesting complementarity between these two aspects of colour constancy, in that simple mechanisms that would tend to maintain excellent colour constancy for one of these types of variation, tend to fail quite badly for the other...

There may not be a general solution that achieves colour constancy with both changing illuminants and changing backgrounds.

In my view, mechanisms that compute cone-excitation ratios are a paradigm example. Whereas illumination changes/variations produce robust invariant signals from such mechanisms, background changes/variations yield widely varying signals, which thus fail to support background-independent constancy. The mechanisms earn their keep by implementing (types of) illumination-independent constancy, not background-independent constancy. Splitting, rather than lumping, thus provides the best framework for investigating such mechanisms.

Another important objection to the chromatic edge hypothesis is that it is motivated by a rare and unusual pathology; indeed, a single subject, whose experience bears little

³⁶ For proposals along these lines, see Maloney (1999: 389) and Brainard & Wandell (1986: 1651).

resemblance to our own. The hypothesis therefore lacks the introspective support and evidence base of mainstream theories of colour constancy. In other work (Davies, 2021), I have presented additional arguments in support of the hypothesis, appealing to aspects of visual phenomenology, physiology, and function in healthy subjects. None of these arguments, however, directly addresses the issue of constancy. I now therefore assess whether it is plausible that edge perception contributes to constancy in healthy subjects.

The empirical evidence is rather equivocal on this matter. Cornellisen and Brenner (1995) conducted an eye tracking study on subjects performing simultaneous asymmetric colour matching tasks. Subjects were given either a paper match or hue-saturation match instruction. If invariant edge appearance contributed to constancy, we would predict a positive correlation between the time spent looking at edges, and the degree of constancy exhibited by subjects' matches. Although subjects had significantly higher degrees of constancy in the paper match condition, however, they spent no more time on average looking at edges than in the hue-saturation match condition (1995: 2440-41). Despite the null finding, it is worth noting that the stimuli were complex Mondrians, in which the test and reference patches abutted numerous differently coloured patches in the background. For the reasons discussed above, this might reduce the effectiveness of using the appearance of the patch's border as a cue in setting one's matches. Another issue highlighted by Foster and colleagues (1997: 1343) is that 'observers may have relied on information derived from edges without looking at them.' It is well established, after all, that subjects can direct attention to features or objects that are not fixated. Absent further data, however, we cannot push this further.

Overall, though, Foster and colleagues (1997) are similarly sceptical regarding the contribution of edges to constancy. Reporting a task involving discrimination of illuminant changes from surface colour changes (1997: 1343),

some control measurements... suggested that edge information has little influence on performance... [O]bservers' ability to discriminate illuminant changes from material changes in Mondrian patterns was found... to be only moderately impaired when black borders of width 1.5 deg visual angle were introduced between patches so that spatial chromatic induction effects were largely eliminated...

Note briefly that the use of Mondrian stimuli complicates matters somewhat, as described above. More importantly, however, Foster and colleagues are not talking about edge perception in the way that I am. Their prediction is that the black border should decrease the extent to which the absolute colour appearance of the patch is influenced by the surround. That is, the border should reduce the effect of 'chromatic induction' or colour contrast, which as we've seen is crucial in maintaining constant colour appearance across changes in illumination. Their point, then, is that insofar as introducing borders only moderately impaired the ability to discriminate illuminant changes from surface colour changes, this task probably does not depend significantly on constancy in the traditional sense of invariant hue and saturation appearance.

There are two important qualifications, however. Firstly, the black borders did at least moderately impair performance. Secondly, note that introducing a large, coincident luminance contrast at the patch's boundary would also have the effect of swamping or masking any contribution of chromatic contrast to the visual appearance of this boundary. That is, the black border might decrease colour contrast, hence hue-saturation invariance; but in addition, it might impair chromatic edge representation, thus precluding invariance in edge appearance. As such, the data are consistent with the view that either invariant hue-saturation or invariant edge appearance contribute at least a small amount to the capacity to discriminate illuminant changes from surface colour changes.

A study by Barbur and colleagues (2004) found more substantial effects of black borders between patches and their surrounds. Their task involved a novel asymmetric matching paradigm, in which the test and reference patch were alternated in rapid succession at the same spatial location. Subjects performed matching by adjustment, using a hue-saturation-brightness match criterion. The task was therefore intended to isolate ‘instantaneous’ mechanisms of colour constancy, in the phenomenal sense of hue-saturation invariance. They reported a ‘rapid decrease’ in constancy with an increase in border width from 0 to 1 degree, levelling off through 3 degrees (2004: 12). Pursuant to the discussion of Foster and colleagues (1997), as this task involved hue-saturation matching, the results are consistent with the hypothesis that the border reduces the effect of colour contrast in stabilising colour appearance under changes in illumination. Again, however, the results are also consistent with the hypothesis that the border undermines the contribution of chromatic edge perception. That is, in some small part, successful hue-saturation matches might involve adjusting the test patch so that its edges appear the same as those of the reference patch in respect of their pronouncedness or distinctness. It must be granted that the first hypothesis carries most weight, but absenting direct evidence to the contrary, the second, edge-based hypothesis remains a live possibility.

In closing, it is worth noting the inherent methodological difficulties in assessing chromatic edge perception in healthy subjects. Whatever the theoretical rationale for positing perceptual representations of edges, it is clear that we ordinarily pay them little heed. Aside from cliffs and cutting objects, edges per se have few practical implications. Barring a few

notable exceptions, edges are not objects of aesthetic appreciation.³⁷ Of course, healthy subjects see colours, and once one has this sort of experience, it is hard, if not impossible, to ignore as a basis for answering colour-related questions. Concerning colour constancy, ‘the edges of that bright red object look constant under the changing light,’ said nobody, ever. Assessing the chromatic edge hypothesis in healthy subjects therefore poses a considerable challenge. That is why the achromatopsic data is, in a sense, a double-edged sword. On the one hand, it may reveal an aspect of constancy that is difficult to discern in normal subjects. On the other hand, for that very reason, we may continue to wonder about its relevance to our own case.

While I feel the force of this concern, I offer two points in mitigation. Firstly, although edges are obscure entities, it seems plausible that perceptually differentiating an object from its background typically requires perceiving at least part of its boundary, as noted on Goldman’s (1977) differentiation condition. As such, if we follow Akins and Hahn in holding that colour vision helps us to see things, then we should find it quite palatable that colour vision helps us discern the edges of objects. By the same token, if the idea that colour vision facilitates object seeing seems relevant to our own case, then the notion that colour vision facilitates the representation of edges should seem very relevant to our own case. Secondly, and in closing, recall that my aim has been to theorise the basic constituents of our mature colour constancy capacity. Given the nature of this project, it is unsurprising that the investigation has yielded a perceptual kind that seems quite unfamiliar, and of uncertain provenance. After all, if it was introspectively obvious that we have chromatic edge

³⁷ Commenting on Josef Albers’s *Aurora*, Murawski (2016) comments that the ‘strong contrasts between hues—as well as alternating squares of light reflection and light absorption—draw our attention to the edges as “hot points” of activity.’

perception, the traditional approach to the philosophy of colour would have singled it out long ago. Having long since left the comfort of the armchair, one can only hope that we now have a clearer view of the lower borders of colour perception.³⁸

Works Cited

Adelson, E. H. & Pentland, A. P. (1996). The perception of shading and reflectance. In D. Knill and W. Richards (Eds.), *Perception as Bayesian Inference*, (pp. 409-423). New York: Cambridge University Press.

Akins, K. (2014). Black and White and Colour. In R. Brown (Ed.), *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*, (pp. 173-224). Springer Press.

Akins, K. & Hahn, M. (2014). More than Mere Colouring: The Role of Spectral Information in Human Vision. *British Journal of Philosophy of Science*, 65(1), 125-171.

Allen, K. (2016). *A Naïve Realist Theory of Colour*. Oxford University Press.

³⁸ An early ancestor of this paper was delivered at the European Society for Philosophy and Psychology, held at the University of Hertfordshire. Many thanks to the audience there, and to my fellow symposiasts Bob Kentridge and Henry Taylor, for valuable discussion and feedback. I am also grateful to audiences at the Universities of Birmingham, Glasgow, and Nottingham, with special thanks to Kathleen Akins, Keith Allen, Derek Brown, Mazviita Chirimuuta, Karl Gegenfurtner, Martin Hahn, Bob Kentridge, Fiona MacPherson, Nick Shea, David Simmons, Maja Spener, and Mark Sprevak. Finally, thanks to Ian Phillips, and to two anonymous referees for helpful comments on drafts of the paper.

- Arend, L. & Reeves, A. (1986). Simultaneous colour constancy. *Journal of the Optical Society of America A*, 3(10), 1743–1751.
- Aristotle. (1991) Sense and sensibilia. In J. Barnes (Ed.) *The Complete Works of Aristotle*, Princeton University Press.
- Barbur, J., Darryl de Cunha, C. W. & Plant, G. (2004). Study of instantaneous color constancy mechanisms in human vision. *Journal of Electronic Imaging*, 13, 15–28.
- Barbur, J., Harlow, A., & Plant, G. (1994). Insights into the different exploits of colour in the visual cortex. *Proceedings of the Royal Society B: Biological Sciences*, 258, 327-334.
- Barbur, J. & Spang, K. (2008). Colour constancy and conscious perception of changes of illuminant. *Neuropsychologia*, 46, 853-863.
- Brainard, D. & Wandell, B. (1986). Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3(10), 1651-1661.
- Brown, D. (2014). Colour layering and colour constancy. *Philosophers' Imprint*, 14(15), 1-31.
- Brown, R. (2003). Backgrounds and illuminants: the yin and yang of colour constancy. In R. Mausfeld & D. Heyer (Eds.) *Colour Perception: Mind and Physical World*, (pp. 247-272), Oxford University Press.
- Burge, T. (2010). *Origins of Objectivity*. New York: Oxford University Press.
- 2014. Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints. *Philosophy and Phenomenological Research*, 88(2), 485-501.

- Byrne, A. & Hilbert, D. (2003). Colour realism and colour science. *Behavioural and Brain Sciences*, 26, 3–64.
- Casati, R. & Varzi, A. (1999). *Parts and Places: The Structures of Spatial Representation*. MIT Press.
- Chirimuuta, M. & Kingdom, F. (2015). The Uses of Colour Vision: Ornamental, Practical, and Theoretical. *Minds and Machines*, 25, 213-229.
- Cohen, J. (2008). Colour constancy as counterfactual. *Australasian Journal of Philosophy*, 86(1), 61-92.
- (2009). *The Red and the Real*, Oxford University Press.
- Conway, B. R. (2001). Spatial Structure of Cone Inputs to Color Cells in Alert Macaque Primary Visual Cortex (V-1). *Journal of Neuroscience*, 21(8), 2768-2783.
- Conway, B. R. & Livingstone, M. S. (2006). Spatial and Temporal Properties of Cone Signals in Alert Macaque Primary Visual Cortex. *The Journal of Neuroscience*, 26(42), 10826–10846.
- Cornellisen, F. & Brenner, E. (1995). Simultaneous Colour Constancy Revisited - an analysis of viewing strategies. *Vision Research*, 35(17), 2431-2448.
- Cowey, A. & Heywood, C. (1995). There's more to colour than meets the eye. *Behavioural Brain Research*, 71, 89-100.
- Craven, B. & Foster, D. (1992). An operational approach to colour constancy. *Vision Research*, 32(7), 1359-1366.
- Davies, W. (2016). Color constancy, illumination, and matching. *Philosophy of Science*, 83(4), 540-62.

——— (2018). Colour vision and seeing colours. *The British Journal for the Philosophy of Science*, 69(3), 657–690.

——— (2021). Colour relations in form. *Philosophy and Phenomenological Research*, 102(3), 574-594.

Foster, D. (2011). Colour Constancy. *Vision Research*, 51, 674-700.

Foster, D., Amano, K., & Nascimento, S. (2000). How temporal cues can aid colour constancy. *Colour research application*, 26(S1), 180-185.

Foster, D., Craven, B., & Sale, E. (1992). Immediate colour constancy. *Ophthalmic and Physiological Optics*, 12, 157-160.

Foster, D. & Nascimento, S. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London B: Biological Sciences*, 257, 115–121.

Foster, D., Nascimento, S., Craven, B., Linnell, K., Cornellisen, F., & Brenner, E. (1997). Four Issues Concerning Colour Constancy and Relational Colour Constancy. *Vision Research*, 37(10), 1341-1345.

Friedman, H., Zhou, H., & von der Heydt, R. (2003). The coding of uniform colour figures in monkey visual cortex. *Journal of Physiology*, 548(2), 593-613.

Gegenfurtner, K. & Kiper, D. (2003). Colour vision. *Annual Review of Neuroscience*, 26, 181-206.

Gert, J. (2010). Colour constancy, complexity, and counterfactual. *Nous*, 44(4), 669-690.

Goldman, A. (1977). Perceptual objects. *Synthese*, 35, 257-284.

Granzier, J. & Valsecchi, M. (2014). Variations in daylight as a contextual cue for estimating season, time of day, and weather conditions. *Journal of Vision*, 14(1), 1-23.

Hardin, C. L. (1988). *Colour for Philosophers*. Indiana: Indianapolis.

Heywood, C., Cowey, A., & Newcombe, F. (1991). Chromatic Discrimination in a Cortically Colour Blind Observer. *European Journal of Neuroscience*, 3, 802-812.

——— (1994). On the role of parvocellular (P) and magnocellular (M) pathways in cerebral achromatopsia. *Brain*, 117, 245-254.

Heywood, C., Kentridge, R., & Cowey, A. (1998). Form and motion from colour in cerebral achromatopsia. *Experimental Brain Research*, 123, 145-153.

Heywood, C., Nicholas, J., & Cowey, A. (1996). Behavioural and electrophysiological chromatic and achromatic contrast sensitivity in an achromatopsic patient. *Journal of neurology, neurosurgery, and psychiatry*, 61, 638-643.

Heywood, C., Wilson, B., & Cowey, A. (1987). A case study of cortical colour 'blindness' with relatively intact achromatic discrimination. *Journal of Neurology, Neurosurgery, and Psychiatry*, 50, 22-29.

Hilbert, D. (2005). Colour constancy and the complexity of colour. *Philosophical Topics*, 33(1), 141-158.

Hurlbert, A. (1996). Colour vision: putting it in context. *Current Biology*, 6(11), 1381-1384.

——— (2003) Colour vision: Primary visual cortex shows its influence. *Current Biology*, 13(7), R270-R272.

- Hurlbert, A., Bramwell, D., Heywood, C., & Cowey, A. (1998). Discrimination of cone contrast changes as evidence for colour constancy in cerebral achromatopsia. *Experimental Brain Research*, 123, 136-144.
- Hurlbert, A. & Wolf, K. (2004). Color contrast: a contributory mechanism to color constancy. *Progress in Brain Research*, 144, 147-160.
- Kentridge, R., Heywood, C., & Cowey, A. (2004). Chromatic edges, surfaces and constancies in cerebral achromatopsia. *Neuropsychologia*, 42, 821-830.
- Koffka, K. (1936). *Principles of Gestalt Psychology*. Kegan Paul, Trench, Trubner & Co. Ltd.
- Kuriki, I. & Uchikawa, K. (1996). Limitations of surface-color and apparent-color constancy. *Journal of the Optical Society of America A*, 13(8), 1622-1636.
- Linnell, K., & Foster, D. (1996). Dependence of relational colour constancy on the extraction of a transient signal. *Perception*, 25, 221–228.
- MacLeod, D. & Boyton, R. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, 69(8), 1183-1186.
- MacPherson, F. (2015). The structure of experience, the nature of the visual, and type 2 blindsight. *Consciousness and Cognition*, 32, 104-128.
- Maloney, L. (1999). Physics-Based Approaches to Modeling Surface Color Perception. In K. Gegenfurtner & L. Sharpe (Eds.) *Colour Vision: From Genes to Perception*, (pp. 387-422), Cambridge University Press.
- Matthen, M. (1999). The disunity of colour. *Philosophical Review*, 108(1), 47-84.

Mollon, J. (1989). “Tho’ she kneel’d in that place where they grew...” The uses and origins of primate colour vision. *Journal of experimental biology*, 146, 21-38.

——— (2006). Monge: The Verriest Lecture, Lyon, July 2005. *Visual Neuroscience*, 23, 297-309.

Mollon, J., Newcombe, F., Polden, P. G., & Ratcliff, R. (1980). On the presence of three cone mechanisms in a case of total achromatopsia. In G. Verriest, (Ed.) *Colour Vision Deficiencies, Vol. 5*, (pp. 130-135), Bristol, England: Hilger.

Murawski, M. (2016). Spotlight Essay: Josef Albers. *Blog of the Kemper Art Museum*, <https://www.kemperartmuseum.wustl.edu/node/11221>.

Nascimento, S., Ferreira, F., & Foster, D. (2002). Statistics of spatial cone excitation ratios in natural scenes. *Journal of the Optical Society of America A – Optics Image Science and Vision*, 19, 1484–1490.

Reeves, A., Amano, K., & Foster, D. (2008). Colour constancy: phenomenal or projective? *Perception and Psychophysics*, 70, 219-228.

Sorabji, R. (1971). Aristotle on Demarcating the Five Senses. *The Philosophical Review*, 80(1), 55-79.

West, M., Spillmann, L., Cavanagh, P., Mollon, J., & Hamlin, S. (1996). Susanne Liebmann in the critical zone. *Perception*, 25, 1451-95.

Westland, S., & Ripamonti, C. (2000). Invariant cone-excitation ratios may predict transparency. *Journal of the Optical Society of America A – Optics Image Science and Vision*, 17, 255–264.

Whittle, P. (2003). Contrast Colors. In R. Mausfeld & D. Heyer (Eds.) *Colour Perception: Mind and Physical World*, (pp. 115-138), Oxford University Press.

Whittle, P. & Challands, P. (1969). The effect of background luminance on the brightness of flashes. *Vision Research*, 9, 1095-1110.

Wright, W. (2013). Colour Constancy Reconsidered. *Acta Analytica*, 28, 435–455.

Zaidi, Q. (1998). Identification of illuminant and object colors: Heuristic-based algorithms. *Journal of the Optical Society of America A – Optics Image Science and Vision*, 15, 1767–1776.

Zaidi, Q., Spehar, B., & DeBonet, J. (1997). Color constancy in variegated scenes: the role of low-level mechanisms in discounting illumination changes. *Journal of the Optical Society of America A*, 14, 2608–2621.