Color Constancy, Illumination, and Matching

Davies, Will

DOI:

[10.1086/687261](https://doi.org/10.1086/687261)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Davies, W 2016, 'Color Constancy, Illumination, and Matching', *Philosophy of Science*, vol. 83, no. 4, pp. 540-562. <https://doi.org/10.1086/687261>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

On an institutional repository or open access repository after 12 months embargo for Non-commercial use. Will Davies, "Color Constancy, Illumination, and Matching," published in *Philosophy of Science* 83, no. 4 (October 2016): 540-562. DOI: 10.1086/687261

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Abstract: Colour constancy is a foundational and yet puzzling phenomenon. Standard appearance invariantism is threatened by the psychophysical matching argument, which is taken to favour variantism. This argument, however, is inconclusive. The data at best support a pluralist view: colour constancy is sometimes variantist, sometimes invariantist. I add another potential explanation of these data, complex invariantism, which adopts an atypical six-dimensional model of colour appearance. Finally I prospect for a unifying conception of constancy among two neglected notions: discriminatory colour constancy and relational colour constancy. The former arguably marks a common core capacity that is present across widely differing viewing contexts.

COLOUR CONSTANCY, ILLUMINATION, AND MATCHING

WILL DAVIES

Imagine that you are looking at a fire engine in a fire station under dim artificial lighting. The vehicle looks red to you – a striking shade of fire engine red, no less. The station alarm sounds, and the engine speeds off down the street. When viewed in natural daylight, the vehicle looks the same vivid shade of colour to you. In general,

Constancy Intuition: A coloured object viewed under different illumination conditions typically looks the same colour across such changes in viewing context.

The Constancy Intuition records one dominant type of response that normal subjects have when viewing coloured objects under varying illumination. These intuitions or naïve judgements of colour constancy are remarkable, given that the proximal stimulation of the visual system notoriously compounds the contribution of stable surface colour and the illuminant. The central aim of the science of constancy is to explain how the visual system pulls off this feat of inversion. What processes and mechanisms are involved in extracting an invariant signal from such variable input?

For all its undoubted progress and ingenuity, there remains a fundamental lack of clarity at the foundations of constancy science. In what sense does a red object ‘look the same colour’ across changes in illumination? The traditional space of options divides into two families of views. The first family holds that colour constancy consists in some invariance in the phenomenology of colour experience, that is, in the qualitative dimensions of conscious *colour appearance*.¹ Let’s call this *phenomenal invariantism*. A competing family holds that colour phenomenology in fact varies significantly across changes in illumination. These are *variantist* views of colour constancy. Variantists differ as to how they explain our constancy intuitions, but historically the most dominant line has appealed to non-phenomenological, quasi-cognitive, processes of perceptual ‘judgement’ or ‘inference’ from variant colour appearance to a stable surface colour property.

This paper focuses critical attention on the *psychophysical matching argument*, which has been taken by vision scientists and philosophers alike to support variantism. The argument draws on some suggestive data from asymmetric colour matching tasks. After introducing the argument in Section 1, I argue in Section 2 that the conclusion of the argument has been overstated. The data at best support a pluralist view, on which colour constancy is sometimes variantist, but sometimes invariantist. In Section 3, I argue that subjects’ viewing strategies vary across differing task conditions, probably resulting in qualitatively different colour experiences. This requires a more nuanced interpretation of the matching data, on which subjects’ matches can be explained to varying degrees by both phenomenal and non-phenomenal factors. In Section 4, I criticise the argument’s assumption

¹ ‘Colour appearance’ is to be understood in the quasi-technical sense familiar to vision scientists, as pertaining to the aspect of perceptual phenomenology represented by multi-dimensional colour appearance spaces (Kuehni 2003; Fairchild 2005).

that the colour appearance of surfaces under varying illumination conditions is exhausted by the traditional three dimensions of hue, saturation, and lightness. I propose that in many contexts, appearance has dimensions of both material and lighting colour, numbering six in total, with colour constancy explained by invariance in the former, despite concurrent variation in the latter. I call this view *complex invariantism*.

Given the evident plurality and complexity of colour constancy phenomena, one might wonder whether there is anything that they share in common. In Section 6, I consider the prospects for a unifying conception by discussing two neglected notions: *discriminatory colour constancy* (DCC) and *relational colour constancy* (RCC). DCC was originally introduced as an operationalization of RCC. Deploying the colour appearance model from Section 4, I first argue that DCC is equally well explained by our capacity to perceptually discriminate material and lighting properties via changes in material and lighting colour appearance. I then argue that DCC is largely phenomenologically neutral. In particular, DCC is consistent with traditional invariantist, traditional variantist, and complex invariantist characterisations of colour constancy. DCC therefore provides a particularly useful orthogonal, cross-cutting, dimension of classification for our colour constancy abilities.

1. The Psychophysical Matching Argument

Phenomenal invariantism holds that colour constancy consists in the stability of colour appearance across changes in illumination. As will be familiar, however, changes in lighting often make an appreciable difference to the visual appearance of objects within the scene. Familiar examples include the appearance of shadows falling across a surface, and the comparatively cool cast of colours viewed under fluorescent bulbs.

Variance Intuition: A coloured object viewed under different lighting conditions appears different in some respect(s) across these conditions.

The variance intuition need not necessarily pose a problem for invariantism. The question is, do these illumination-dependent variations in appearance consist in changes in *colour* appearance? And if so, are these variations in colour appearance significant enough to undermine the invariantist's proposed phenomenal explanation of colour constancy?

Variantists provide positive answers to both of these questions. In paradigmatic instances of colour constancy, they claim, colour appearance exhibits a high degree of illumination-dependent variation:

Variant Colour Appearance: In standard cases, an object's colour appearance varies significantly and systematically with changes in the illumination.

Variant Colour Appearance implies that colour appearances typically are not remotely stable under lighting changes. If this is right, then our commonplace intuitions of colour constancy will not be explicable via any phenomenal stability across changes in viewing context:

Non-Phenomenal Colour Constancy: Our colour constancy intuitions cannot be explained by an invariance in the qualitative dimensions of colour appearance.

Taken together, these two claims spell trouble for phenomenal invariantism.

The strongest grounds for Variant Colour Appearance and Non-Phenomenal Colour Constancy derive from matching arguments, which come in two forms: intuitive and psychophysical. Noë (2004) presents an intuitive matching argument. Consider a variably illuminated object such as a white wall partly in shadow. According to Noë, we have a robust constancy intuition regarding the colour of the wall: it seems to us that we perceive the wall, in some sense, as uniform in colour. Now suppose you are presented with a range of colour samples, and asked to select the samples that best match a) the directly illuminated region of the wall, and b) the shaded region of the wall. Noë (2004, 128) argues that different colour samples would be selected in each case, and that this implies a difference in colour appearance between the two regions. The inference is that if these regions are matched by different colour samples, then the difference in phenomenal character caused by the difference in illumination must itself be a difference in colour appearance, as per Variant Colour Appearance.

The psychophysical matching argument brings some experimental rigour to these claims. The argument first appeared in the philosophical literature in Thompson (1995, 197), and is developed by Cohen (2008).² It centres on some famous experiments by Arend and Reeves (1986), employing an asymmetric matching paradigm. In this paradigm, subjects are presented with two stimuli: the ‘standard field’, with a uniform reflectance R presented under illuminant I_1 , and the ‘test field’, with the same reflectance R presented under a different illuminant I_2 . The subject is instructed to adjust the *chromaticity* of the test field until she achieves a match between the test field and the standard field. The chromaticity of a stimulus signifies its coordinates in the CIE chromaticity diagram.³ The degree of colour constancy exhibited can be represented using a Colour Constancy Index (CCI), with values ranging from 1 (idealised or perfect constancy) to 0 (absence of constancy).⁴

Intuitively, the CCI value represents the degree to which the difference in illumination between the two fields affects the response of the subject’s colour perception system to these stimuli. The higher the CCI, the smaller the implied effect of the illumination change on her perceptual state, and hence the better her implied colour constancy.

The crucial twist in Arend and Reeves’ experiments is that prior to making any matches, subjects received one of two possible task instructions. The first was to ‘match the hue and saturation of the test field to that of the standard field’. The labels ‘hue’ and ‘saturation’ are taken to refer to dimensions of the hue, saturation, and lightness colour appearance space. Call this the *appearance match condition* (AMC). The second was to ‘make the test field look as though it is cut from the same piece of paper’. Call this the *surface match condition* (SMC). Subjects in the two conditions performed significantly differently on the matching task. Subjects in the AMC achieved fairly low CCIs, averaging 0.2. In contrast, subjects in the SMC achieved much higher CCIs, averaging 0.52. In a

² The argument is endorsed by Brad Thompson (2006, 80ff) and Millar (2013, 222, fn. 8), and receives brief critical attention in Jagnow (2009, 570) and Hilbert (2005, 157, fn. 7). See also Wright (2013, 438ff).

³ The chromaticity diagram is derived from the CIE 1931 2° Standard Observer, in which the axes represent the proportions of carefully selected lights or ‘primaries’ that are required to match the stimulus in respect of colour. See Fairchild (2005, 77-78) and MacAdam (1985) for details.

⁴ See Arend et al. (1991).

replication of these results, Bäuml (1999, 1537-1541) found average CCIs of 0.23 in the AMC, and 0.79 in the SMC.⁵

Arend and Reeves' results provide satisfying confirmation of the competing intuitions of constancy and variance, familiar to us all when viewing an object under different illuminants. A plausible interpretation of the differing CCIs is that subjects in the AMC are attending to some relatively variant aspect of their colour perception system's overall response to the stimuli, whereas subjects in the SMC are attending to some reasonably invariant aspect of their perceptual response. Given that subjects in the AMC had been instructed to attend to the hue and saturation of the stimuli, a natural inference is that these appearance dimensions underwent significant illumination-dependent variation, as per Variant Colour Appearance. Cohen (2008, 67-68) argues as follows,

When subjects make appearance matches... they make the regions cease to be discriminable (along whatever dimension they were previously discriminable) by adjusting the hue and saturation of one of them. Now, it is a standard assumption in visual psychophysics that the hue and saturation of a patch are dimensions of its apparent colour; if so, then adjusting the hue and saturation of the test patch just is adjusting the patch's apparent colour. Therefore, whatever the difference was in virtue of which the patches were initially visually discriminable, that difference can be offset by a difference in apparent colour. And this, in turn, might lead us to suspect that the difference revealed in the [variance intuition] is a difference in apparent colour...

To summarise, the positive conclusion of the argument is captured by Variant Colour Appearance: when viewing stimuli of equivalent reflectance under different illuminants, the phenomenal character of subjects' colour experience exhibits significant illumination-dependent variation. The associated negative conclusion is captured by Non-Phenomenal Colour Constancy: our colour constancy intuitions therefore cannot be explained by an invariance in the qualitative dimensions of colour appearance. What *does* explain our constancy intuitions on this view? This is all up for grabs: variantists differ as to how they ultimately explain our capacity for colour constancy. The force of the argument is simply to leave phenomenal invariantism moribund, with some or other form of variantism our best remaining option.

2. What is the Scope of the Psychophysical Matching Argument?

The psychophysical matching argument has been taken by its philosophical advocates to warrant quite sweeping conclusions about the nature of colour constancy. I read Cohen (2008), for example, as arguing for variantism about colour constancy simpliciter, across the board and without qualification. This fits a clear pattern within the philosophy of perception,

⁵ These transformed CCIs are provided by Foster (2011, 683). Across 13 comparable simultaneous asymmetric matching studies, average appearance match CCIs ranged from 0.11-0.46, and surface match CCIs from 0.35-0.86.

in tacitly assuming that colour constancy is a unified phenomenon, admitting of a single theoretical treatment. On broader inspection of the empirical literature, however, this uniqueness assumption seems frankly untenable.

As an example, one basic distinction in the study of colour constancy is between *simultaneous* and *successive* constancy. Simultaneous colour constancy reflects our ability accurately to judge, within a scene at a time t , that two objects under different illuminants have the same surface colour. Successive constancy reflects our ability accurately to judge that a single object under varying illuminations conditions at times t and $t+n$ remains the same colour. The focal study by Arend and Reeves concerns simultaneous colour constancy. Arend and Reeves' stated aim (1986, 1743) was to isolate the contribution of 'simultaneous mechanisms' of constancy. They accordingly presented the standard and test fields concurrently, side by side on a monitor. Cohen (2008, 63) nonetheless thinks that a similar pattern of results is to be found in successive constancy,

[I]n cases of successive colour constancy we find the same pair of perceptual reactions that occur in cases of simultaneous colour constancy. On the one hand, normally sighted subjects find that the two (successively presented) regions of interest are, in some sense to be explained, alike in apparent colour. And on the other hand, normally sighted subjects find that the two (successively presented) regions of interest are, in some sense to be explained, easily, obviously, and quickly visually discriminable in apparent colour.

It seems clear, then, that Variant Colour Appearance is taken to apply to both simultaneous and successive forms of colour constancy.

Cohen is surely right that in many cases, successive illumination changes will produce strong variance intuitions. This fits our experience, for example, of a lamp being flicked on and off, or a cloud's shadow passing over a green field. In other cases, however, things are less clear cut. While rapid changes in illumination caused by moving shadows and such like are extremely common, an even more regular feature of our environment is the slow, gradual, change in natural light that occurs over the course of the day. Now crucially, it is well established that changes in illumination that occur over longer durations allow greater scope for adaptation by the visual system; and adaptation to the illuminant is known to have a *significant* normalising effect on colour appearance.⁶ Indeed Arend and Reeves (1986, 1743) raised precisely this point, noting that 'there are extensive data showing that adaptation, alone or in combination with simultaneous mechanisms, can produce large hue and saturation shifts.' This gives us reason to doubt that successive colour matches involving longer timescales will follow the pattern of Arend and Reeves' results. This is confirmed by Kuriki and Uchikawa (1996), who administered a dichoptic successive matching task in which subjects were preadapted to the illuminant for fifteen minutes. They found CCIs of 0.72 in the AMC, and 0.77 in the SMC. This all but eliminates any difference in performance

⁶ Cornelissen and Brenner (1991). By 'adaptation' I mean *light* adaptation. So-called *contrast* adaptation also plays a significant role in colour constancy. See Webster and Mollon (1995) and Brown and MacLeod (1997).

between these two task conditions, with subjects displaying comparably high CCIs in both conditions. These data do not fit the pattern of Variant Colour Appearance. In fact, in these viewing conditions, the data are consistent with, and adequately explained by, phenomenal invariantism.⁷

The conclusions of the psychophysical matching argument therefore need to be hedged appropriately. Properly qualified, the positive conclusion should be that colour appearance exhibits a significant degree of illumination-dependent variation *under some conditions*, such as conditions of low adaptation. Similarly, the negative conclusion should be that *in such conditions*, our colour constancy intuitions are not explained by an invariance in the qualitative dimensions of colour appearance. This suggests a more complex picture of colour constancy than is typically assumed in the philosophical literature. In summary,

Pluralism: There exist many different types of colour constancy, with differing perceptual natures, which will be given differing psychological explanations.

The idea behind Pluralism is that our naïve intuitions of colour constancy – our judgements that things in some sense look the same colour across changes in illumination – will be explained by different underlying perceptual capacities in different contexts. These explanations will differ in whether they appeal primarily to phenomenological or non-phenomenological factors, for example, or whether they fit invariantist or variantist specifications. A moderate pluralism was in fact proposed by Arend and Reeves (1986, 1743, 1749), and later reiterated by Reeves and colleagues (2008, 220). In a detailed investigation of this literature, Wright (2013) endorses a very similar view. In what follows, I shall be developing and expanding on this pluralist theme in several ways, yielding an even more complex and heterogeneous picture.

3. Viewing Strategies and the Unfixed Nature of Phenomenology Across Task Conditions

The psychophysical matching argument has both a positive and negative conclusion. The purported evidence for the positive conclusion, Variant Colour Appearance, is fairly well understood: subjects in the AMC are asked to match stimuli in respect of the appearance dimensions of hue and saturation, and they achieve very low CCIs. The purported evidence for the negative conclusion, in contrast, is far less clear. What exactly is the basis for Non-Phenomenal Colour Constancy?

Let's start by considering subjects' performance in the SMC. In this condition, subjects are instructed to make the stimuli look as though they are cut from the same piece of paper. Subjects here achieve comparably much higher CCIs than in the AMC. As such, we can infer that whatever aspect of their perceptual response subjects are attending to in the SMC, this aspect is apt to explain colour constancy, such as it is, in such conditions. This invites the question as to what this constancy-grounding perceptual aspect *is* exactly. As noted above, variantists differ in their views on this matter. The key negative claim is simply that, whatever else it might be, this aspect does not include any invariance in respect of colour appearance. How does this follow? The variantist's reasoning, it seems, is that we

⁷ I have gained much from Wright's (2013) discussion of these issues.

know from the AMC data that subjects have highly illumination-variant colour experiences when attempting appearance matches with exactly the same stimuli. Assuming that subjects undergo similarly illumination-variant experiences of these stimuli when in the SMC, it is reasonable to expect that their colour experiences similarly exhibit insufficient invariance to explain their higher CCIs. In summary, the phenomenology implied by the matching data from the AMC is assumed to be fixed across task conditions, in the following sense. For subject *X* who achieves a CCI of, say, 0.25 with respect to stimuli *A* and *B* in the AMC, her colour experience of *A* and *B* in the SMC will exhibit a sufficiently similar degree of illumination-dependent variation to that in the AMC, such that any increase in CCI in the SMC is not attributable to an increase in appearance invariance in the SMC. Let's call this the *fixed phenomenology assumption*.

I want to apply some moderate pressure to this assumption. I argue that subjects' colour experiences in the SMC might in fact undergo a partial change from the AMC, as a result of the change in task conditions. Although changes of this particular sort would be insufficient to explain the entire difference in CCIs, a relatively more invariant colour appearance in the SMC might provide part of the explanation of the higher CCIs recorded in this condition. The proposed mechanism is that the change in task instructions influences subjects' *viewing strategies*, which in turn influences the degree of illumination-invariance in their colour experiences. By a 'viewing strategy', I mean a combination of factors including the direction and distribution of gaze across the stimuli; the amount of attention assigned to various parts or features of the stimuli; and the amount of time spent viewing parts of the stimuli. My argument draws on a study by Cornelissen and Brenner (1995). Their hypothesis is that the differing task instructions could alter the amount of time that subjects spend looking at certain parts of the stimuli, with the consequence that subjects in the SMC are more adapted to the illuminant than in the AMC. Such differences in adaptation in turn could alter the character of subjects' colour experiences, which might partly explain the differences in their CCIs.

Cornelissen and Brenner's findings require careful interpretation. They did observe a significant change in subjects' viewing strategies between the two task conditions. In particular, subjects spent longer looking at the stimulus surround before making a match in the SMC than in the AMC (1995, 2437). They then estimated the likely state of adaptation for each subject, given the amount of time exposed to the stimulus surround, and calculated the expected influence of this adaptational state on their colour experience. This predicted adaptational effect was then compared against the size of the difference in CCIs between the two task conditions. The results differed significantly from subject to subject. Two subjects exhibited much larger differences in CCI than predicted by their likely difference in adaptation. These subjects happened to have the most marked difference in CCI between the two conditions. It was predicted that differences in colour experience explained at best between 10-15% of this difference in constancy performance. For three further subjects, the effects of instruction were consistent with an explanation largely due to eye movements and adaptation. Complicating these results, however, is the fact that these subjects did not produce marked differences in matches across the AMC and SMC (1995, 2435).

What are we to conclude from these data? For subjects exhibiting a marked difference in CCIs between task conditions, changes in viewing strategy and subsequent adaptation clearly do not explain all of this difference. Nonetheless, as Cornelissen and Brenner (1995, 2439) note, such changes 'probably do contribute to some small extent to colour constancy

[in the SMC].’ Contrary to the fixed phenomenology assumption, then, it is a mistake to assimilate wholesale the implied appearance data from the AMC with matching data from the SMC. The AMC data tell us about the degree of illumination-dependent variance exhibited by subjects’ colour experiences when asked to match stimuli for hue and saturation, but need not accurately reflect the degree of variance encountered by subjects even in closely related task conditions.⁸

Contrary to Non-Phenomenal Colour Constancy, it is thus premature to dismiss appearance invariance as completely irrelevant to subjects’ improved CCIs in the SMC. Subjects in the SMC are likely to undergo colour experiences with slightly higher degrees of invariance than in the AMC. Whereas variantists would like to attribute subjects’ colour constant surface matches entirely to some non-phenomenal process, the evidence suggests that moderate improvements in appearance invariance could in fact play some role. The moral to draw is that the seemingly binary distinction between phenomenal and non-phenomenal colour constancy presents a false dichotomy. While Pluralism claimed that colour constancy can be sometimes phenomenal, sometimes non-phenomenal, it is perhaps more accurate to say that constancy can be more or less phenomenal, with appearance invariance explaining subjects’ colour constant matches to a greater or lesser extent, depending on context. In summary,

Complexity: For many surface matches, a subject’s CCI will be explained by a combination of phenomenal factors, such as appearance invariance, and non-phenomenal factors, which might include processes of perceptual ‘inference’. The extent to which each type of factor is explanatorily relevant will vary with viewing context.

4. The Nature of Colour Appearance

Following a widely accepted standard in colour science, Arend and Reeves (1986) and those following them have assumed a three dimensional, hue, saturation, and lightness model of colour appearance. Three dimensional models of colour date back at least to the 18th Century. Significant theoretical advances were made by figures such as Munsell at the turn of the 20th Century. The three dimensional Munsell Colour System of hue, value (i.e. lightness), and chroma (or ‘purity’, related to saturation) remains one of the most common tools for representing colour appearance.⁹ From the 1950s, more rigorous studies of colour appearance were enabled by multidimensional scaling (MDS) techniques, in which subjects are asked to compare colour stimuli for pairwise (dis)similarity, and the results analysed in order to extract dimensionality and distances between stimuli in a metric space. MDS studies have been taken to provide some of the strongest confirmations of the three dimensional nature of colour appearance. For example, in seminal work Indow (1988) performed 19 MDS experiments, which recovered a very close approximation of Munsell’s original three dimensional space.

⁸ Cornelissen and Brenner only assessed one possible mechanism by which subjects’ colour experiences might vary across task conditions, namely adaptation. This leaves it open as to whether other mechanisms might further stabilise colour appearance in the SMC.

⁹ Kuehni (2003) provides a fantastic historical overview. See also Fairchild (2005, 96-99).

Given this history, it is predictable and to some extent reasonable that colour appearance is so routinely assumed to have three dimensions. In the present context, however, this assumption is highly questionable. It is an under-appreciated fact that the aforementioned models of colour appearance were generated using colour stimuli presented under neutral and homogenous illumination. The same applies to Indow's key MDS studies. This invites the question as to whether three dimensions of colour appearance are adequate to represent stimuli viewed under varying illuminants. As we'll now see, if this assumption fails, the invariantist is presented with a number of additional responses to the psychophysical matching argument.

One such response is due to Hilbert (2005, 152). Hilbert endorses the idea that when viewing variably illuminated stimuli, colour appearance has additional dimensions that represent the way that an object is illuminated. While Hilbert does not recommend any specific model, he cites as an example the five dimensions proposed by Fairchild (2005, 91ff): brightness, lightness, colourfulness, chroma, and hue. Hilbert's strategy is to appeal to such additional dimensions to defuse Variant Colour Appearance. He concedes that subjects' matching performance in the AMC is explained by illumination-dependent variations in colour appearance, but argues that these variations occur in dimensions such as brightness and colourfulness, rather than hue and saturation/chroma. Hilbert thus retains the core invariantist view that constancy consists in the stability of hue and saturation, which he takes to represent objects' surface spectral reflectance properties.

Hilbert's response faces a number of problems. As Wright (2013, 449) has argued, whatever the merits of supplementing the traditional three dimensions of colour appearance, it is unmotivated to claim that subjects in the AMC are attending to anything other than hue and saturation. This is because these subjects were explicitly instructed to attend to these dimensions. One could argue that subjects might have been inaccurate or fallible in following these task instructions. Perhaps subjects simply misunderstood, or made systematic errors in assessing and reporting their own phenomenology. While clearly possible, however, such scepticism is largely muted by the fact that subjects in Reeves and colleagues' (2008) replication study were given significant training in performing hue and saturation matches. These subjects performed the task with confidence and consistency, making the suggestion of misunderstanding or systematic error seem ad hoc.

I propose an alternative strategy. Whereas Hilbert sought to explain away Variant Colour Appearance, let's grant that subjects in the AMC are following task instructions and performing hue and saturation matches. As in the previous Section, I want to focus attention on Non-Phenomenal Colour Constancy. Another assumption that drives the inference from Variant Colour Appearance to Non-Phenomenal Colour Constancy is that single dimensions of hue and saturation (and lightness) *exhaust* the possible dimensions of colour appearance to which subjects could attend in either task condition:

Exhaustiveness: If subjects' performance in the AMC is explained by their attending to single dimensions of hue and saturation (and lightness), then no further dimensions of colour appearance are available to provide an explanation for subjects' performance in the SMC.

The introduction of additional dimensions of colour appearance allows us scope to deny Exhaustiveness. I argue that subjects in the SMC are also tracking phenomenal dimensions of

hue and saturation, albeit hue/saturation dimensions of a different kind to those tracked in the AMC. This requires distinguishing two different types of hue and saturation: *material* or *surface* hue and saturation, on the one hand, and *lighting* or *illumination* hue and saturation, on the other. Let's call this the *duplicate dimensions gambit*. This gambit allows for a phenomenal explanation of both the appearance matching and surface matching data. The view ultimately will be that subjects in the SMC are matching stimuli by attending to material dimensions of colour appearance, while subjects in the AMC are matching via lighting dimensions of colour appearance.

The duplicate dimensions gambit has a long history within colour science, although as Mausfeld (2003, 386) notes the approach has remained well outside the mainstream. In important recent work on this theme, Tokunaga and Logvinenko (2010a, 2010b) sought to reassess the dimensionality of colour appearance using MDS methods, but allowing variations in illumination across the colour stimuli. Tokunaga and Logvinenko found that in these conditions, a minimum of six appearance dimensions were required to model subjects' dissimilarity judgements. This is broadly in keeping with previous duplicate dimension views, on which our experience of a uniform surface under coloured light involves two *distinct* phenomenal attributes: a material colour and a lighting colour. Total colour appearance is specified by two independent triples of hue, saturation, and lightness.¹⁰ In a break from such views, however, Tokunaga and Logvinenko argue that their six colour dimensions are not best modelled as two independent attributes. They claim that 'both triplets constituting these six numbers are determined by a pair object/light, not object separately and light separately' (2010c, 2556).¹¹ On their analysis, then, the dimensions of material and lighting colour are modelled as constituents of a single, complex, object colour attribute. Our visual experience of object hue, for example, is thus 'essentially 2D' (2010c, 2556): we do not experience material blue or lighting yellow simpliciter, for example, but rather material blue under a yellowish illuminant.¹²

Relating back to our central puzzle, this view yields a second-generation invariantism, on which colour constancy consists in the approximate invariance of material colour appearance, while allowing for significant variation in lighting colour appearance. Let's call this *complex invariantism*, as distinguished from *traditional* invariantist models incorporating just three colour dimensions.¹³ The conjecture is that subjects' responses in the AMC are driven either primarily by lighting hue and saturation, or some complex function of lighting colour and material colour; whereas subjects in the SMC are attending primarily to material hue and saturation. Complex invariantism thereby grants a version of Variant Colour Appearance while resisting Non-Phenomenal Colour Constancy, providing a satisfying synthesis of traditional invariantist and variantist views.

5. The Limits of Complex Invariantism

¹⁰ Mausfeld (2003, 388) describes these as 'two different representational primitives.' See also MacLeod (2003).

¹¹ See Tokunaga and Logvinenko (2010b, 1744).

¹² See Matthen (2010) for a similar view.

¹³ Compare Tokunaga and Logvinenko (2010c, 2555-2556). Note that traditional variantist views likewise presuppose that colour appearance has just three dimensions.

Complex invariantism is not intended as a decisive response to the psychophysical matching argument. My aim was to cast doubt on the move from Variant Colour Appearance to Non-Phenomenal Colour Constancy, by questioning the assumption of Exhaustiveness.

Exhaustiveness is far from secure, given recent evidence that colour appearance has multiple dimensions of hue and saturation in conditions of multiple or varying illuminants. Given that Arend and Reeves' experiments involve simultaneously viewing two stimuli of equivalent reflectance under differing illuminants, we have reason to think that these conditions would involve such complementary dimensions. If that's right, then complex invariantism should provide a good alternative explanation of the matching data. This conjecture is tentative, however, and must await empirical confirmation. In keeping with Pluralism, moreover, we should note that in scenes with single or constant illumination, the traditional three dimensions of colour appearance might suffice to model subjects' dissimilarity judgements (Logvinenko and Tokunaga 2011, 444). Traditional forms of invariantism and variantism may therefore still be required to explain our constancy intuitions in such contexts, such as they may be.

A further qualification is that the mere presence of dimensions of material and lighting colour in a context does not guarantee that colour constancy will be best explained by invariant material colour appearance. We should allow that in some contexts, both lighting colour *and* material colour appearance could exhibit significant illumination-dependent variation. These dual colour spaces therefore open the door to complex forms of variantism, as well as complex invariantism. Traditional variantists hold that subjects 'infer' an object's surface colour from patterns of illumination-dependent variation in hue, saturation, and lightness. A complex variantist, in contrast, could hold that subjects appeal to variations in dimensions of both lighting colour and material colour, in inferring a stable surface colour. Development of this idea must await another occasion.

A third issue concerns our interpretation of Arend and Reeves' task instructions in the AMC. In this condition, subjects are instructed to 'match stimuli in respect of hue and saturation'. One interpretation of complex invariantism noted above is that subjects subsequently match stimuli in respect of lighting hue and saturation. This would raise the question, however, as to why subjects disambiguate the task instructions this way. Why default to matching lighting hue and saturation, rather than material hue and saturation? One explanation is that subjects are cued to adopt this reading by those administering the task. In the training regime of Reeves and colleagues (2008), for example, perhaps subjects were *taught* that 'hue and saturation' matches should be performed by equating stimuli for lighting hue and saturation. The nagging follow-up question, of course, is why those administering the tasks default to this understanding of 'hue' and 'saturation'. I do not know the answer to this question, but speculate that the history of these technical terms is so closely bound up with colorimetric studies of the colours of lights, that those inculcated into existing scientific practice acquire recognitional concepts of 'hue' and 'saturation' that denote lighting dimensions of colour appearance. This issue is clearly contentious, however, and requires more detailed consideration than I can provide here.

A fourth issue concerns the interpretation of subjects' dissimilarity judgements in Tokunaga and Logvinenko's studies. Subjects were presented with two coloured papers, each under different lighting, and instructed to 'evaluate dissimilarities between papers'. Tokunaga and Logvinenko deliberately avoided using terms such as 'colour appearance', and 'colour stimulus' in their instructions. Moreover they 'did not specify explicitly the colour

dimensions between which the dissimilarity was supposed to be measured. The intention was to ascertain the dimensions which will emerge from the multidimensional analysis of dissimilarities rather than to impose some dimensions on observers.’ (2010b, 1741) Given that the aim of any such MDS study is to generate a model of subjects’ colour appearance space, then, Tokunaga and Logvinenko reasonably judge that they should not presuppose a specific view of appearance in their task instructions.¹⁴

While such neutrality is therefore required by the nature of the MDS study, it presents a fundamental problem. Given the indeterminacy in the instruction to ‘evaluate dissimilarities between papers’, how can we be sure that subjects are specifically and exclusively assessing features of colour appearance? Can we rule out, for instance, that subjects were assessing overall dissimilarity in part by appeal to colour appearance, but also in part by appeal to the output of some non-phenomenal or quasi-cognitive process? Variantists may well object that it is question-begging to rule out such alternative explanations of the data. After all, the very issue at hand concerns the appropriate interpretation of subjects’ competing intuitions of similarity and dissimilarity when viewing coloured stimuli under differing illuminants. The variantist’s proposal is that our intuitions of dissimilarity reflect the illumination-dependent nature of colour appearance, whereas our intuitions of similarity reflect the outputs of some non-phenomenal process. And perhaps this variantist model could explain why six dimensions are required to model subjects’ dissimilarity judgements: three dimensions to model variant colour appearance, as standard, and a further three dimensions to model the content of the variantist’s putative non-phenomenal process.

This issue threatens a serious deadlock. On the one hand, complex invariantism appeals to the MDS data to vindicate an appearance-based explanation of colour constancy. On the other hand, variantists may seek to fit the data to their own preferred non-phenomenal account of constancy. It is not altogether clear who has the upper hand here. One point that tells against the variantist response, however, is that it requires some selectiveness in interpreting MDS studies. As with almost everyone, variantists are happy to adopt the three dimensional models of colour appearance generated by previous MDS studies, in which illumination conditions were kept uniform. Given that Tokunaga and Logvinenko’s studies adopt the same general approach, it seems unprincipled to explain away half of their data as irrelevant to colour appearance. Indeed the variantist seems even guiltier of begging the question here. The variantist seeks to explain half of the dissimilarity data by appealing to their preferred non-phenomenal account of colour constancy. But as argued above, Non-Phenomenal Colour Constancy is only motivated given the assumption of Exhaustiveness, which in turn presupposes that colour appearance has only three dimensions. Given that the argument for Non-Phenomenal Colour Constancy contains this presupposition, it is most definitely question-begging to appeal to the argument’s conclusion to fend off challenges to the three dimensional model of colour appearance. Of course, there may well be other, more independent, routes to blocking these challenges. But as things stand, the onus is on the variantist to provide such independent grounds.

6. Discriminatory and Relational Colour Constancy

¹⁴ This point indeed applies to any MDS study of this nature.

Let's take stock. Colour constancy is a highly heterogeneous and complex phenomenon. The phenomenological basis for colour constancy should be assessed on a case-by-case basis, with appearance invariance – either in the traditional three appearance dimensions, or in material colour appearance – playing more or less of an explanatory role, depending on context. The psychophysical matching argument therefore misses its mark: there is no straightforward dialectical route from asymmetric matching to a generic variantism about colour constancy. The evidence suggests a far more nuanced and diverse picture.

I now want to complicate matters more by introducing two further concepts of colour constancy: *discriminatory colour constancy* (DCC) and *relational colour constancy* (RCC). The seminal presentation of these notions is due to Craven and Foster (1992). While this work has received a fair amount of attention in vision science, it has been largely neglected in philosophy. This is unfortunate, for Craven and Foster (1992, 1359) are motivated by the very puzzle under discussion here, concerning the appropriate response to Arend and Reeves' famous results. After noting the difficulties posed by these results for 'the traditional definition of colour constancy' (1992, 1360), i.e. phenomenal invariantism, they present

an alternative and complementary property of colour constancy...: The ability of a subject to correctly attribute changes in the colour appearance of a scene either to changes in the spectral composition of the illuminant or to changes in the reflecting properties of that scene. This aspect of colour constancy is not concerned with the nature or extent of any changes in colour appearance, but simply with the subject's interpretation of them. For a certain area of a scene to be identified by a subject as physically unchanged under a change in illuminant, it is not necessary that the area generate some invariant local percept; all that is needed is that it is perceived to stand in relation to other areas in the scene in the same way.

This passage introduces two related capacities. The first is a *discriminatory* capacity: the ability accurately to distinguish changes in surface spectral reflectance from changes in illumination. I refer to this capacity as DCC. The second is a *relational* capacity, on which colours are somehow 'perceived to stand' in constant relations of similarity and difference across changes in illumination. This capacity is RCC, on which constancy consists not in the stability of monadic colour appearance, but in the (approximate) invariance of perceived colour relations across changes in context. My questions are as follows. What exactly is the connection between DCC and RCC? And what is the relationship between these notions and the views of colour constancy discussed earlier?

In the remainder of the paper, I argue for two claims. First, although Craven and Foster originally presented DCC as an operationalization of RCC, RCC is not the only possible basis for DCC. I suggest that these discriminatory capacities are better explained in some contexts by the subject's awareness of changes in the complementary dimensions of material or lighting colour appearance. Second, I propose that DCC provides a dimension of classification for our colour constancy capacities that is orthogonal to the invariantist-variantist characterisations discussed above. The concept is strictly neutral as to whether the phenomenology of colour constancy is traditional invariantist, traditional variantist, or complex invariantist. As such, DCC provides a useful unifying characterisation of colour

constancy, a kind of lowest common denominator view of the constancy-related abilities involved across widely differing viewing contexts.

DCC was presented by Craven and Foster as an alternative to phenomenal invariantism, jettisoning the requirement that colour appearance should be stable across changes in illumination. DCC requires merely that subjects are able to discriminate changes in illumination from changes in surface material properties. What explains our capacity to make these discriminations? What perceptual phenomena are involved? According to Craven and Foster, 'all that is needed' to perform such discriminations is for surfaces viewed under differing illuminants to be perceived, in some sense, as standing in the same relations of similarity and dissimilarity. In other words, the proposed perceptual basis for DCC is RCC. As later explained by Foster and colleagues (1997, 1342), 'discriminating illuminant changes from material changes... corresponds to discriminating whether the relations between surface colours are unchanged.' Perceptual awareness of these colour relations thus 'underlie[s] the ability of observers to discriminate, reliably and effortlessly, illuminant changes from material changes in scenes' (1997, 1341).

How exactly does RCC explain DCC? The idea seems to be as follows. RCC with respect to stimuli *A* and *B* involves the perceptual awareness of the invariant colour relations between *A* and *B* across changes in illumination. By the same token, it is assumed that subjects will be aware of some variance in the colour relations between *A* and *B* if the surface colour of either stimulus is changed. DCC is grounded in a) the subject's capacity to discriminate between changes in colour appearance that preserve colour relations and those that don't, and b) the ability to judge as follows: given some change in distal conditions *C*, and subsequent transformation *T* in the colour appearance of *A* and *B*, if *T* preserves the colour relations between *A* and *B*, then *C* is likely to have been an illuminant change; if not, *C* is more likely to have been a surface material change.

I do not have the space here to assess whether this explanation of DCC is plausible. I will note, however, that RCC remains an extremely underdeveloped notion, and would benefit greatly from more sustained philosophical attention. In particular, it is unclear in what sense objects are 'perceived to stand' in invariant or variant colour relations. Craven and Foster never adequately explained this claim, and their brief attempts at pinpointing the phenomenology of RCC are quite confusing.¹⁵ I will discuss these issues in more detail in future work. In any event, RCC does not provide the only possible explanation of DCC. Indeed in contexts of multiple or varying illuminants, it may not even provide the best explanation. Drawing on the model developed in Section 4, a much simpler and clearer explanation of DCC in such contexts appeals to the phenomenal dimensions of material and lighting colour appearance. The proposal is satisfyingly straightforward: our ability to discriminate a material change from a lighting change reflects our ability to distinguish a change in material colour appearance from a change in lighting colour appearance. In such contexts, DCC is explained by the subject's sensitivity to the phenomenological difference between changes in material colour appearance, as contrasted with changes in lighting colour appearance.

¹⁵ See in particular Craven and Foster (1992, 1364) and Foster and colleagues (2001, 287-288).

It is interesting to contrast this proposal with Craven and Foster's view. On their view, DCC provides an operationalization of RCC. DCC is grounded in the subject's capacity to perceptually discriminate whether a transformation in (three dimensional) colour appearance preserves or violates the colour relations between objects in the scene. Subjects then judge on this basis whether the transformation is likely to have been caused by an illuminant change or a material change. On their view, then, subjects do not strictly *perceptually* discriminate changes in material or lighting properties. Such changes are merely 'attributed' (1992, 1360) on the basis of the perceived colour relations in the scene. On my view, in contrast, our ability to discriminate illuminant changes from material changes is grounded in the subject's capacity to perceptually discriminate changes in material properties and lighting properties themselves, via awareness of changes in material and lighting colour appearance. These two explanations have quite different implications for the perceptual representations involved in DCC. On my view, DCC is explained by perceptual capacities that function to represent monadic properties of both surface material and lighting in the scene. On Craven and Foster's view, DCC is explained by appeal to the subject's perceptual awareness of colour relations. As noted above, it remains unclear how we should characterise this putative relational awareness, and what kinds of perceptual representation it might involve. As Craven and Foster intended it, however, it seems fairly clear that RCC does not involve representations of surface material and lighting properties.

To recap, I've argued that in some contexts, DCC will be well explained by the six-dimensional colour appearance model developed in Section 4. In keeping with the spirit of Pluralism and Complexity, however, I suggest that DCC will receive different explanations in different viewing contexts. Assuming the details can be worked out, for instance, Craven and Foster's explanation of DCC in terms of RCC may prove the most plausible in contexts in which colour experience has the traditional three dimensions of variation. As a general concept of colour constancy, however, DCC is satisfyingly broad: it can be applied in widely differing viewing contexts, in which the phenomenology of colour constancy can vary significantly. Indeed I shall now argue that DCC is strictly neutral as to whether the phenomenology of colour constancy fits the specifications of traditional invariantism, traditional variantism, or complex invariantism.¹⁶

From the preceding paragraph, let's assume that DCC can be explained in a context either by RCC, or by the subject's awareness of changes in the phenomenological dimensions of material and lighting colour. From the latter case, DCC is clearly consistent with a complex invariantist account of the phenomenology of colour constancy. What about the former case? I argue that when DCC is explained by RCC, it is consistent with either a traditional invariantist or traditional variantist view of colour constancy. Consider two coloured stimuli *A* and *B* presented under illuminant L_1 , where *A* appears red and *B* appears orange. Grant that *A* appears redder and darker than *B*. Let *A* and *B* viewed under L_1 be mapped to locations (w, x) and (y, z) in the chromaticity diagram. Now suppose that *A* and *B* are presented under a different illuminant L_2 . Let the subsequent change in the appearance of *A* and *B* be modelled by the transformation of chromaticity coordinates from (w, x) to $(w',$

¹⁶ This claim might at first seem surprising, given that Craven and Foster (1992, 1360) introduced DCC as an alternative to traditional phenomenal invariantism. As we'll see, however, although DCC does not *entail* invariantism, it is *consistent* with it – as it is with variantism.

x'), and (y, z) to (y', z') . The idea behind RCC is that colour constancy is consistent with any such change in appearance, so long as the illumination transformation preserves all colour relations between A and B . For example, the location (w', x') should represent a colour that is both redder and darker than that represented by (y', z') . We should note, however, that this constraint is trivially satisfied in the null case in which the transformation maps (w, x) and (y, z) onto themselves. That is to say, it is trivial that if A and B appear exactly the same colour across the illumination change, then A and B will appear to stand in the same colour relations.¹⁷ Formally speaking, then, RCC is consistent with both traditional invariantist and variantist accounts of the colour appearance of A and B across changes in illumination. As such, insofar as DCC is to be explained by RCC, DCC is consistent with both traditional invariantist and variantist accounts of colour constancy. Combining these results, DCC is consistent with traditional invariantist, traditional variantist, and complex invariantist accounts of constancy. As a concept of colour constancy, then, DCC does not imply any one view of the phenomenology of colour constancy: DCC, let's say, is *phenomenologically neutral*.

One might object that where the colour appearance of A and B is completely invariant constant across changes in illumination, our explanation of DCC *need not* appeal to RCC. In such cases, subjects could judge that there has been no change in surface material – and hence that there must have been a change in illumination – simply by introspecting the invariant monadic colour appearance of A and B . My view is that it is an empirical question as to what explains subjects' discriminatory capacities in such contexts. My prediction would be that, as elsewhere, subjects' judgements probably will be explained by multiple factors: in this case, perhaps appearance invariance, in part, and RCC, in part. In any event, the objection does not undermine the point that *insofar as* RCC explains DCC, DCC is strictly speaking neutral as between traditional invariantism and variantism.

DCC's phenomenological neutrality makes it extremely appealing in the present context. In Sections 2 and 3, I argued that the extant matching data do not decisively favour traditional variantism over traditional invariantism. The principles of Pluralism and Complexity suggest a far messier and context-bound view of the phenomenological bases for colour constancy. This picture was muddled further by the introduction of complex invariantism in Section 4. This heterogeneity is at once both illuminating and frustrating. On the one hand, our taxonomic powers with respect to subjects' constancy abilities hopefully have been expanded and refined. On the other hand, it is natural to despair at the fragmentation of colour constancy into such a ragbag of phenomenological types. Would it be too much to hope for a disambiguation of 'colour constancy' that picks out some more general, fundamental, psychological kind?¹⁸ DCC's phenomenological neutrality presents us with a unifying opportunity here. Where there is colour constancy of any type, I propose, there will be discriminatory colour constancy. One thing that even the most extreme invariantist and variantist types of constancy will have in common, I suggest, is that the perceptual capacities involved will support a high degree of DCC. This proposal won't satisfy those who have set their lofty sights on theorising *the* phenomenology of colour constancy. But if the foregoing view is right, there's no such unique thing.

¹⁷ The possibility of invariantist RCC is discussed by Foster and Nascimento (1994, 119).

¹⁸ Somewhat related concerns are expressed by Foster (2003) under the heading 'Does Colour Constancy Exist?'

7. Conclusion

In closing, I want to address an issue that has been kept firmly in the background, concerning the relationship between colour constancy and colour ontology. For a time, there was hope that colour constancy could provide clues as to the nature of the colours. In its traditional invariantist guise, colour constancy was taken by some to support colour physicalism, the view that colours are surface spectral reflectances.¹⁹ The underlying assumption was that if colour appearances are largely illumination-independent, then so must be the colours. As Cohen (2008) has argued, if phenomenal invariantism is an inaccurate characterisation of colour constancy, then this undercuts the claimed support for colour physicalism. By the same token, if traditional variantism is true, this seemingly favours the view that colours are constituted by illumination-dependent properties. This approach appears fundamentally misguided, however, in light of the extreme heterogeneity and complexity of colour constancy. There is no straightforward, context-free, answer as to whether colour appearance is illumination-invariant or not. Indeed the very interpretation of 'colour appearance' likely differs from one context to another. A proper understanding of colour constancy therefore won't help settle these longstanding debates about ontology. In emphasising colour constancy's pluralist and complex nature, however, the foregoing discussion may help reroute colour ontology into less doggedly sectarian territory.

¹⁹ Byrne and Hilbert (2003).

References

- Arend, L., and Reeves, A. 1986. "Simultaneous colour constancy." *Journal of the Optical Society of America A* 3 (10): 1743–51.
- Arend, L., Reeves, A., Schirillo, J., and Goldstein, R. 1991. "Simultaneous colour constancy: papers with diverse Munsell values." *Journal of the Optical Society of America* 8:661–72.
- Bäumel, K. H. 1999. "Simultaneous Colour Constancy: How Surface Colour Perception Varies with the Illuminant." *Vision Research* 39:1531–50.
- Brown, R. O., and MacLeod, D. I. A. 1997. "Colour appearance depends on the variance of surround colours." *Current Biology* 7:844–49.
- Byrne, A., and Hilbert, D. 2003. "Colour realism and colour science." *Behavioural and Brain Sciences* 26:3–64.
- Cohen, J. 2008. "Colour constancy as counterfactual." *Australasian Journal of Philosophy* 86 (1): 61–92.
- Cornelissen, F., and Brenner, E. 1991. "On the role and nature of adaptation in chromatic induction". In *Channels in the Visual Nervous System: Neurophysiology, Psychophysics, and Models*, ed. Blum, B. 109–23. London/Tel Aviv: Freund.
- 1995. "Simultaneous Colour Constancy Revisited: an Analysis of Viewing Strategies." *Vision Research* 35 (17): 2431–48.
- Craven, J., and Foster, D. H. 1992. "An operational approach to colour constancy." *Vision Research* 32 (7): 1359–66.
- Fairchild, M. D. 2005. *Colour Appearance Models, 2nd Edition*. Reading, MA: Addison-Wesley.
- Foster, D. 2003. "Does Colour Constancy Exist?" *Trends in Cognitive Sciences* 7:439–43.
- 2011. "Colour Constancy." *Vision Research* 51:674–700.
- Foster, D., Amano, K., and Nascimento, S. 2001. "Colour constancy from temporal cues: Better matches with less variability under fast illuminant changes." *Vision Research* 41:285–93.
- Foster, D., and Nascimento, S. 1994. "Relational colour constancy from invariant cone-excitation ratios." *Proceedings of the Royal Society, London, B* 257:115–21.
- Foster, D., Nascimento, S., Craven, B., Linnell, K., Cornelissen, F., and Brenner, E. 1997. "Four Issues Concerning Colour Constancy and Relational Colour Constancy." *Vision Research* 37 (10): 1341–45.
- Hilbert, D. R. 2005. "Colour constancy and the complexity of colour." *Philosophical Topics* 33 (1): 141–58.

Indow, T. 1988. "Multidimensional studies of Munsell colour solid." *Psychological Review* 95:456–70.

Jagnow, R. 2009. "How representationalism can account for the phenomenal significance of illumination." *Phenomenology and the Cognitive Sciences* 8 (4): 551-72.

Kuehni, R. G. 2003. *Colour Space and its Divisions: Colour order from antiquity to present*. New Jersey: John Wiley and Sons.

Kuriki, I., and Uchikawa, K. 1996. "Limitations of surface colour and apparent colour constancy." *Journal of the Optical Society of America A* 13:1622–36.

Logvinenko, A. D., and Tokunaga, R. 2011. "Colour Constancy as Measured by Least Dissimilar Matching." *Seeing and Perceiving* 24:407-52.

MacAdam, D. L. 1985. "The physical basis of colour specification", In *Readings on Colour. Vol. 2: The Science of Colour*, eds. Byrne, A., and D. Hilbert, 33-63. Cambridge, MA: MIT Press.

MacLeod, D. I. A. 2003. "New dimensions in colour perception." *Trends in Cognitive Science* 7:97–99.

Matthen, M. 2010. "How things look (and what things look that way)." In *Perceiving the World*, ed. Nanay, B., 226-253. New York: Oxford University Press.

Mausfeld, R. 2003. "'Colour' as part of the format of different perceptual primitives: the dual coding of colour." In *Colour Perception: Mind and Physical World*, eds. Mausfeld, R., and D. Heyer, 381-430. New York: Oxford University Press.

Millar, B. 2013. "Colour constancy and Fregean representationalism." *Philosophical Studies* 164:219-31.

Noë, A. 2004. *Action in Perception*. Cambridge, MA: MIT Press.

Reeves, A., Amano, K., and Foster, D. 2008. "Colour constancy: phenomenal or projective?" *Perception and Psychophysics* 70:219-28.

Thompson, B. 2006. "Colour Constancy and Russellian Representationalism." *Australasian Journal of Philosophy* 84 (1): 75-94.

Thompson, E. 1995. *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. London and New York: Routledge.

Tokunaga, R., and Logvinenko, A. D. 2010a. "Material and lighting hues of object colour." *Ophthalmic and Physiological Optics* 30:611–17

——— 2010b. "Material and lighting dimensions of object colour." *Vision Research* 50:1740–47.

——— 2010c. "Hue manifold." *Journal of the Optical Society of America A* 27 (12): 2551-57.

Webster, M. A., and Mollon, J. D. 1995. "Colour constancy influenced by contrast adaptation." *Nature* 373:694–98.

Wright, W. 2013. "Colour Constancy Reconsidered." *Acta Analytica* 28:435–55.