Colour layering and colour relationalism[1]

Derek H. Brown (brandon University, brownd@brandonu.ca)

Abstract. Colour Relationalism asserts that colours are non-intrinsic or inherently relational properties of objects, properties that depend not only on a target object but in addition on some relation(s) that object bears to other objects. The most powerful argument for Relationalism (Cohen 2009) infers the inherently relational character of colour from cases in which one's experience of a colour contextually depends on one's experience of other colours. Experienced colour layering – say looking at grass through a tinted window and experiencing opaque green through transparent grey – demands a contextual interdependency of one's experience of one of these colours on one's experience of the other. However, most if not all colour ontologies, and core perceptual experiential mechanisms like acquaintance and representation, can accommodate colour layering. It follows that experienced colour layering is consistent with colours being non-relational – this contextual interdependency of colours does not entail the constitutive dependency of one colour on the other. I seek to bring out this tension between a well-known argument for Relationalism and colour layering, and argue that our justification for Relationalism is far weaker than Relationalists suggest. I first introduce readers to colour layering (§1), then to Relationalism (§2), and following this focus on the intersection of these topics (§3).

§1 A primer on colour layering

The aim of this section is to briefly overview how colour layering connects to experience, ontology, and colour constancy. This is essential to bolster the impact layering has on Relationalism (§3).

Pretheoretically, we think of and talk about transparent things or objects as existing and often having their own colours: liquids (beer), solids (glass, plastic), and gases (mist).[2] We arguably draw at least some of

---

[1] For the final version of this paper please see:
Colour layering and colour relationalism. *Mind and Machines: A Special Issue on the Philosophy of Colour* ed. by M. Chirimuuta, *25(2):* 177-191.
For helpful comments on earlier drafts of this paper, many thanks to Jonathan Allan, Mazviita Chirimuuta, Jonathan Cohen, Wayne Wu, an anonymous referee at *Minds and Machines*, and audiences at the Pacific APA 2015, the University of Manitoba, and Towards a Science of Consciousness 2015.

[2] Setting aside colours for a moment, here are rough experiential and causal characterizations of 'transparent', 'transluscent' and 'opaque'. Experientially, to be transparent is to be something that can be clearly seen through, causally it is to transmit (as opposed to reflect) light without distortion. An object is transluscent to the extent that things are seen less clearly through it, and as it begins to distort (e.g., scatter) light that it transmits. Opaque objects cannot be seen through, and transmit no light. In this context 'thing' or 'object' must apply at least to liquids, solids and gases.

these thoughts and utterances from colour experience. Pretheoretically, there is no difficulty with experienced colour layering, that is, with experiencing two colours along a line of sight, one opaque colour through a transparent one. When one looks at grass through the tinted window of one's car, by hypothesis one's experience can be of opaque green through transparent grey. In this way there is a *contextual interdependency* of one's experience of one of these colours on one's experience of the other. This is consistent with each colour, and one's experience of each colour, being distinct – this interdependency does not entail conflation.

Some rough distinctions help sketch the space of possible experiences one might have when looking in a direction that contains (for simplicity) two objects, a proximal and distal object:

(1) *Colour layering:* when one experiences distinct colours belonging to each object along the line of sight.

(2) *Colour fusion:* when one experiences a single colour that mixes or fuses together the contributions of the objects located along that line of sight.

(3) *Single, non-fused colour:* when one experiences a single colour that only reflects the contributions of one of the things located along that line of sight. Here the other object makes no contribution to one's colour experience – it is in this way invisible. This can occur for example because the proximal object is completely invisible to the subject (e.g., a perfectly transparent sheet of glass whose edges are hidden from view), or because the proximal object is occluding the distal one.

My primary interest is in colour layering and its distinctness from colour fusion (I will largely set (3) aside), concepts and a distinctness I take to be pretheoretically straightforward and uncontroversial.

When we shift to colour theory, matters need be no different. I briefly embed colour layering within different colour ontologies and connect the discussion to colour constancy.[3] The injection of colour fusion into the remainder of this section is not difficult and will be left to the reader.

---

[3] In psychology, layered visual experiences – *scissions* – have been of particular interest since the 1970s and are actively studied today. Metelli (1974) is a classic work. Kitaoka (2005) corrects some deficiencies, and Anderson (2008) offers an alternative and summarizes some seemingly fatal problems for Metelli's view and the views based on it. Kingdom (2011) is an excellent review article of relevant phenomena, and Casati (2009) provides an good summary and connects the issue to some philosophical debates. See also, e.g., Khang & Zaidi (2002a,b) and Wollschläger and Anderson (2009). It is fair to say that there is no general account of scissions that dominates contemporary empirical literature. Instead, we should expect that scissions will remain an active area of empirical research at least in the near-term.

Theories of colour ontology can accept transparent colours and colour layering, though naturally differ on how to analyze it. Objectivists like Byrne and Hilbert (2003) argue for a generalized notion of colour *productance* that spans not only surface colours, but also the colours of filters, volumes and light sources. Although they don't discuss contexts involving potential colour layering in detail, their ontology supports its occurrence. Mentalists – those holding that colours are purely properties of minds – can accept that there are transparent coloured things and colour layering. They will deny that any objective things have colours and attribute colours instead to whatever mental entities they take to bear them. For example, if sense-data are mental objects and the bearers of colours, then one should admit the possible existence of transparent and opaque coloured sense-data, and of contexts in which they co-occur and thus are layered. Eliminativists – deniers of any instantiators of colour in our world – can allow for the possible (i.e., instantiated in some world) existence of transparent colours and colour layering, and for example utilize representation to explain our seeming experiences of such things in our world. Even relationalists, broadly conceived, can countenance transparent colours and colour layering, provided that all colours are defined relative to something. Thus if one experiences green grass through a grey window, Relationalists can hold that the greenness of the grass is not intrinsic to the grass but is a relational property of the grass plus other factors, and so also for the greyness of the window. I will return to the relationship between Relationalism and colour layering in §3. The general point is that no colour ontology need antecedently fear transparent colours or colour layering.

The same can be said for broad approaches to colour experience. Most generally, the two core mechanisms for perceptual experience are representation and acquaintance. Roughly, the acquaintance mechanism explains experience solely by reference to relations between perceivers and perceptual objects, be those objects objective things or subjective entities like sense-data. By contrast representation explains experience most centrally by appeal to *contents* (e.g., uninstantiated universals). I take it as adequately straightforward that colour layering can be expressed via either mechanism and also by hybrid approaches that utilize both mechanisms. There are thus no antecedent conflicts between colour layering and influential approaches to perceptual experience.

Colour layering directly connects to *colour constancy*, one of the more difficult issues in colour theory. Roughly speaking, colour constancy involves a stability in the experienced colour of a thing despite the presence of a variation in some perceptual factor such as illumination conditions. For example, when one experiences a partly shadowed object or an object in daylight then under dense cloud, *and* in some sense experiences that object to have a constant colour, then one undergoes a colour constancy experience. There is thus both a *constant* and a *variable* element in colour constancy scenarios, though matters become controversial when one attempts to provide detail. Note that although discussions of constancy typically emphasize illumination conditions as the variable element, the phenomenon is more general than this. Nothing precludes the variable element from being a transparent filter or volume. Further, nothing precludes the variable element from being surface colour and the constant element from being a filter, volume or condition of illumination. Here the "proximal" colour remains constant while the "distal" colours do not. For example we can experience a room as similarly illuminated despite the presence of a host of surface colours before us. This is a kind of experienced colour constancy.[4]

The constancy cases directly relevant to layering are cases in which the variable and constant element are both (in some sense) experienced.[5] Here a means of explaining what is experientially constant that incorporates what is variable into our experiences must be sought, and layering – e.g., experiencing a constant element by looking at it through a variable element, or vice versa – has natural application (see xxx). Suppose the elements are a transparent film in front of an opaque surface. Assume that one experiences colour layering when looking at the scene. If one of these colours remains constant, while the other varies, and we experience precisely that, then one experiences a constant colour despite the simultaneous presence and experience of a variable colour. That is, colour constancy occurs, and it is explained by reference to experienced colour layering. It is (at least logically) straightforward to extend this layering analysis to cases in

[4] This suggestion goes against the spirit of some constancy models, where preference is given to the distal or opaque surface colours over the contributions filters, volumes and illuminants make to visual stimuli. However, this tension is not a legitimate philosophical reason to exclude this kind of case from colour constancy.

[5] These cases are to some extent in conflict with a traditional computational approach to colour constancy which hypothesizes that the variable element is *discounted* by our visual systems, and so not experienced by us, or at minimum that this is the relevant *goal* or *ideal* of our visual systems. Land & McCann (1971) and Land (1986) are well-known examples, and Brainard & Wandell (1988), Wandell (1989) Brainard, and Brunt & Speigle (1997) offer different approaches. Jameson & Hurvich (1989) is one well-known response to "discount the illuminant" approaches. Shevell & Kingdom (2008) is a worthwhile recent review of the wider literature.

which a volume is an element instead of a filter, and to cases in which illumination is an element instead of a filter or volume, and of course to cases involving various combinations of these elements. There are many worries one might have about the widespread application of a layering analysis of constancy. In my view such worries are generally misplaced, but detailing why goes outside the scope of this work. For dialectical purposes I ask the reader to take for granted the tenability of a layering analysis of numerous types of constancy cases, including those involving experienced illuminations and surfaces.

§2 A Primer on Colour Relationalism

Colour Relationalism proper is the broad thesis that a thing's colour is not intrinsic to it, but only persists by virtue of relations that thing bears to other things. This alone says nothing about what class of objects (O_c) instantiate these relational colours, nor about the class of relations (R_c) that constitute colours. Regarding the former, for example: O_c may only contain objective things in the world, like dogs and trees; O_c may only contain mental objects like sense-data; O_c may only contain objects that do not exist in our world but do exist in some possible set of worlds; O_c may contain some combination of these or other collections of objects. Similarly, there are many options for the class of relations that constitute colours, including: R_c might only contain relations between objective things; R_c might only contain relations between mental objects; R_c might contain relations between both objective and mental objects.

I will not give names to each of these variants of Relationalism, or explore them in any detail. My point is twofold. First, it is important to emphasize that Relationalism can be articulated within colour mentalism, dispositionalism, objectivism and eliminativism. The above characterization reinforces this. Second, evidence offered in favour of Relationalism must be assessed against options of the above sort to determine the extent to which that evidence in fact supports Relationalism, and if so the extent to which it supports some variant(s) of Relationalism over others. As one might expect, this turns out to be a complicated matter, and it is not one that I will delve into. My interest is instead in a general form of reasoning that has been offered in support of Relationalism and its connection to colour layering. I proceed with the former.

One well-known source of evidence for Relationalism is simultaneous contrast effects involving focal objects and the visible objects and colours above, below and to the right and left of the object.[6] Call this collection of relations the *surround* of the object and this the class of *center-surround* instances of the relativity of colour. Note that by stipulation an object's surround excludes what is in front of and behind the object. The experienced colour of an object can change dramatically when it is surrounded by one set of experienced colours and then another. For example an object can be experienced as blue when embedded in one surround and green when embedded in another. Notoriously, there are surround (or contrast) colours – for example pure black, pure white, brown and olive green – that is, colours that are only experienced with suitable surround contrasts in place. Center-surround contrast phenomena are not only dramatic but now essential to our understanding of colour experience.

The temptation to infer Relationalism is natural.[7] The experienced colour of an object *A* fundamentally depends on the experienced colours of objects $B_{1\text{-}n}$ that surround it (this is an instance of the *Relativity of Colour Perception)*. While one might try to single out a preferred surround for experiencing *A*'s colour, after inspection there is often no principled, well-motivated reason for doing so (*No Privilege*). Epistemically, we thus seem obligated to give these different colour experiences of *A* equal weight (*Ecumenicism*). Object *A* seems intrinsically the same across these variations in its surrounds (*Object Stability*). Therefore, we have grounds to infer (e.g., the best explanation of the above is) not merely that *A*'s *experienced* colour depends on the experienced colours of objects surrounding *A*, but that *A*'s colour *proper* depends on the colours proper of objects surrounding it. In such a case I will say that Relationalism has *Explanatory Import*.

There are various means of resisting this reasoning. One could seek to find a privileged – e.g., "normal" or "ideal" – context in which to perceive *A*'s colour. One could deny *Object Stability*, perhaps on grounds that *A* is a mental object (e.g., a sense-datum) that does intrinsically change across these kinds of contexts. One could generally resist any attempt to infer a categorical conclusion about colour from data drawn from colour experience. Thus, it may be that, because of our visual systems, for us to experience

[6] In this context the relations above, below, to the right and to the left of are defined relative to the perspective of the perceiver. Thus, what is to the right of an object for one perceiver might be to the left of an object to another perceiver.
[7] What follows roughly accords with Cohen's Master Argument (2009, p. 24). Space prevents a thorough comparison.

brown suitable contrast must be present, but this does not entail that brown cannot be instantiated without suitable contrast present (e.g., Byrne & Hilbert 2003). This has the potential result of colours being very difficult to learn about via colour experience (*Epistemicism*).

These disputes are not my concern. For simplicity I will assume that the lesson to draw from center-surround contrast effects is that some variant(s) of Relationalism are true. Which ones? Here are some initial thoughts. If, as is unlikely, the correct explanation of center-surround contrast locates the reason for the effects in objective relations between objective things in our world, then an objectivist Relationalism receives support from the phenomenon, and subjectivist ones do not. If, as is most likely, the reason pertains to peculiarities of our visual system, then some form of subjectivist Relationalism is mandated to the exclusion of objectivist ones. But to this point there is no means of distinguishing, for example, between a subjectivist Relationalism in which the members of O_c are objective things and some members of R_c are subject-involving (e.g., dispositionalism), and a subjectivist Relationalism in which the members of O_c are subjective things like sense-data and the members of R_c are subject-involving. Regarding the latter, the sense-datum theorist can hold that brown sense-data only occur when appropriate contrasting sense-data do too. Thus, center-surround contrast data are at least at first pass equally-consistent with dispositionalist and mentalist forms of Relationalism. Considerably more options are available when one delves with the Eliminativist into possible worlds (which I will not do).

Center-surround contrast effects afford a reasonable quantity of evidence for Relationalism, since these effects occur throughout our colour-visual lives. However, this class of effects is but one of the many classes that fall under 'the relativity of colour'. Others include variations due to successive contrast effects, to illumination effects, to additional peculiarities of perceivers' perceptual apparati (e.g., types of cones), et cetera. If there was reason to believe that the above reasoning should be applied to all or at least most instances of the relativity of colour, then we would have a general and powerful argument for Relationalism. Such a generalization of reasoning would follow from a *Uniformity Hypothesis* to the effect all or most instances of the relativity of colour should be treated as above, namely that each case can be described in a parallel fashion, and that there are grounds to apply *No Privilege, Ecumenicism, Object Stability* and therefore *Explanatory*

Import. This is what Cohen maintains. At the beginning of *The Red and the Real* he nicely states the explanatory import of Relationalism for instances of center-surround simultaneous contrast. When concluding the discussion he states:

> I'll be emphasizing that the very same considerations can be applied again and again across different sorts of variations…I take this fact to be significant, since I think the recurrence of *identical considerations* across different sorts of variation puts pressure on us to favor a uniform reaction in each case – a lesson that, as we'll see, many philosophers have failed to observe. (2009, 20-21, italics added).

In many respects this remark sets the tone for the book, and I find no wavering from it in the work.

It is worth emphasizing that with the *Uniformity Hypothesis* Relationalism transitions from an ontological constraint helpful to explain some set of phenomena to *the* constraint essential to explaining a vast collection of phenomena that is among the most challenging for colour theorists. With the *Uniformity Hypothesis* Relationalism effectively becomes a default view.

What I wish to challenge via colour layering is this *Uniformity Hypothesis*. The point is *not* to argue that there is no evidence for Relationalism – I have, for the purpose of this discussion, conceded that there is. The point is also *not* to argue that Relationalism is at odds with some set of data – on this I wish to remain mute. The point is instead to argue that it is remarkably premature to postulate the *Uniformity Hypothesis*, and therefore that we lack a motive to regard Relationalism as the explanatory saviour Cohen presents it as.

§3 Colour layering and Relationalism

For simplicity I will split the relativity of colour into three classes. The classes are not meant to be exhaustive, nor need each class be exclusive on further examination. The point is to find a means of isolating the impact of colour layering on the justification for Relationalism. Here is the proposal, noting that terminology will be explicated below:

(1) *Center-surround relativities:* instances in which the experienced colour of a thing varies as experienced objects above, below, to the left or right of the object vary.

(2) *Line-of-sight (LOS) relativities:* instances in which the colour experience along a LOS varies as experienced objects in front of or behind an object vary.

(3) *Perceiver-object relativities:* instances in which the experienced colour of a thing varies as unexperienced aspects of perceivers' perceptual apparati vary.

Center-surround instances of simultaneous contrast fall under (1). By hypothesis variations in experienced colour due to numbers of cone types and cone sensitivities fall under (3), as do variations due to differences in other aspects of the architecture of perceptual systems. There is a long list of phenomena falling under (2):

> *LOS relativities* can include colour perceptions through: sunglasses, windows and other transparent films; liquids, atmospheres and other transparent volumes; plausibly the vitreous humour and lenses of one's eyes; the light illuminating a scene; afterimages and other line of sight instances of successive contrast; et cetera.

It is this class that is relevant to assessing the impact colour layering has on the justification for relationalism.

My argument against the *Uniformity Hypothesis* has two steps. The first is the conceptual claim that experienced colour layering undermines the reasoning that leads to *Explanatory Import*. The second is the empirical claim that any presumption of the widespread absence of colour layering requires an at present unmotivated bias, for example a bias in favour of Relationalism. I sustain this not by demonstrating the proliferation of colour layering in experience, but by arguing that there may be widespread layering throughout experience, and that there are several factors that make it difficult to assess the extent to which layering in fact occurs. The prudent position is therefore some form of agnosticism and hence not the *Uniformity Hypothesis*.

§3.1 Layering and the explanatory import of relationalism.

In §1 I noted that experienced colour layering – say looking at grass through a tinted window and experiencing opaque green through transparent grey – demands a *contextual interdependency* of one's experience of one of these colours on one's experience of the other. One experiences the transparent colour by looking through it to the opaque colour, and vice versa. In §1 we also saw that most if not all colour ontologies, and core experiential mechanisms like acquaintance and representation, can accommodate colour layering. Hence

experienced layering is consistent with each colour being non-relational – this contextual interdependency of colours does not entail the constitutive dependency of one colour on the other.

The basic means by which layering targets the reasoning for *Explanatory Import* is by offering a scenario in which two colours are uncontroversially intertwined in a context, both ontologically and experientially, but in which there is (I argue) no temptation to infer that, because of this relation, those colours are *constitutively* dependent on one another. By contrast the reasoning for *Explanatory Import* roughly recognizes the dependency in a perceptual context of an experienced target colour on another colour (e.g., the dependency of an experienced center colour on an experienced surround colour), and infers from this that the nature of the target colour constitutively depends on the other colour (i.e., infers Relationalism). In short one infers a constitutive dependency from a contextual one. While this inference, though not without its critics, is fairly compelling when applied to simultaneous contrast cases, cases involving colour layering illustrate a novel respect in which the inference is questionable.

Let me explain in more detail by focusing on the following scenario[8]:

> *Layering Scenario.* Suppose there is an opaque white object O_{white}, a transparent grey object T_{grey}, and a transparent yellow object T_{yellow}. At first you are presented with O_{white} behind T_{grey}, and then with O_{white} behind T_{yellow}. Suppose that in both contexts you experience colour layering, so that at first you colour experience O_{white}-*through*-T_{grey} and then O_{white}-*through*-T_{yellow}.

Consider two potential objections to this description of the scenario.

One might ask whether describing your colour experiences via iterations of the complex O_x-*through*-T_y can be simplified, say to merely describing your experiences via iterations of O_z. That is, aren't you just experiencing the opaque thing to be differently coloured across these contexts? This is not a friendly amendment, but instead an attempt to eliminate experienced colour layering. There is no such simplification: experienced colour layering can *only* be adequately described via a complex such as O_x-*through*-T_y. In a similar vein one might ask: Does the colour of *O* look different across these contexts? In some sense yes and in

[8] For the purposes of this section it is helpful to ignore the impact of non-LOS factors on colour experience.

some sense no. I suggest that in this discussion focusing on "looks" locutions, well-known troublemakers in perceptual theory, is liable to mislead our analysis. Let me therefore attempt to keep us on track.

What is important about this Scenario is that we may additionally suppose that the object colours are constitutively non-relational, that our three objects are intrinsically white, grey and yellow, respectively. This is reasonable because if these objects do have intrinsic colours, and you are first presented with O_{white} behind T_{grey}, and then with O_{white} behind T_{yellow}, then, if you experience colour layering, we would expect you to first experience O_{white}-*through*-T_{grey} and then O_{white}-*through*-T_{yellow}. This is to say that there is absolutely no pressure to infer from the contextual interdependence of one's experiences of these colours to the constitutive dependence of one of these colours on another. There is not only no pressure to infer Relationalism from this scenario, doing so would be illegitimate.

For concreteness, consider how the Scenario impacts the reasoning for *Explanatory Import*. *Object Stability* is reinforced, since putting T_{grey} and then T_{yellow} in front of O_{white} suggests no intrinsic change to any of these objects. *No Privilege* and *Ecumenicism* are also reinforced, since there is no reason to treat either the colour experience O_{white}-*through*-T_{grey} or O_{white}-*through*-T_{yellow} as somehow privileged over the other, and hence both should be treated equally. And the Scenario is clearly an instance of the relativity of colour perception. So why does the explanatory import of Relationalism not follow?

The reason is because the argument for *Explanatory Import* demands a specific interpretation of the relativity of perception, and that interpretation does not apply in layered scenarios. Here is Cohen's description of the relativity of colour perception: "There are multiple, psychophysically distinguishable perceptual effects (in respect of color) of a single color stimulus" (2009, 24). I will take his 'effects' to be interchangeable with my 'experiences'. Clearly O_{white} is causally involved in bringing about the experience/[effect] O_{white}-*through*-T_{grey}, and in that sense O_{white} has a role in generating these distinct colour experiences in the Scenario. But it would be erroneous to hold that O_{white} itself generates these distinct colour experiences, or to hold that these distinct experiences are simply experiences of the colour of O_{white}. The latter in particular is why the reasoning for *Explanatory Import* does not go through. In order for *No Privilege, Ecumenicism,* and *Object Stability* to create a rationale for Relationalism, the variable colour experiences drawn

from the relativity of colour must be experiences of the target object's colour (i.e., "of a single color stimulus"). But in our Layering Scenario the difference between the experiences O_{white}-*through*-T_{grey} and O_{white}-*through*-T_{yellow} is in a fundamental sense not a difference in the experienced colour of O_{white}, it is a difference in the experienced colour of something else along the LOS on which O_{white} is currently located. Thus while, as it were, the total colour experience along this LOS changes, and O_{white} is the terminal visible object along that LOS, the experienced colour of O_{white} is stable, and there is no motive to infer that O_{white}'s colour is relational.

This is not to say that layered contexts and LOS relativities more generally can all be treated similarly. Most importantly, this class of relativities contains instances of experiential colour fusion. Recall that during fusion one is presented with multiple objects along a LOS, but one experiences a single colour that mixes or fuses together the contributions of the objects located along that LOS. If I put a blue sheet of cellophane a few centimeters in front of a yellow book, I (by hypothesis) can experience O_{yellow}-*through*-T_{blue}. However, if I put the cellophane directly on the book (by hypothesis) the only colour I experience is green. In the latter case the experienced colour is a fusion of the colour contributions made by the cellophane and book. Experienced colour layering is *not* an adequate tool for describing such a case. By contrast, hypothesizing that the experienced green colour is inherently relational – that it depends on colour contributions from two distinct objects, and on those objects standing in a particular relation to one another – is, though not forced, at first pass credible. However, it is worth gaining some clarity on the significance of this by returning to the reasoning for *Explanatory Import*.

In the above fusion case *Object Stability* is reinforced, neither the book nor the cellophane are inherently changed in the scenario. However, the stabilities of the "experiential objects" are arguably not preserved. When there are a few centimeters between the cellophane and the book, both are experienced as distinct objects, and the distinctness of their colours are (by hypothesis) experienced. However, when the distance between the two object bounds toward zero, their experienced colours are fused together. Perhaps this is in part because the objects are no longer experienced as distinct: one doesn't experience a sheet of cellophane on a book, but a single object with a single facing surface to which green is attributed. If so then there is a single experienced colour that is attributed to some single experiential object. One could debate

about how this connects to Cohen's interpretation of the relativity of colour, but I wish to set that aside. What is important is the impact this fusion scenario has on our analysis of *No Privilege*.

If one's aim is to discern the colour of the book, then a colour experience in which the book's contribution to colour is distinctly experienced should be preferred over one in which the book's and cellophane's contributions are fused together. Fusion experiences yield information, not about the book's colour *simpliciter*, but about a relation between the book's colour and the cellophane's colour. This relational information is valuable, but if one is seeking information about the book's colour, the information is arguably inferior to information in which the book's contribution to colour experience is distinctly packaged, or separated out from the contributions of other objects. The contrast with layered experiences is instructive: the layered experience O_{yellow}-*through*-T_{blue} distinguishes the book's contribution to the experience from the cellophane's; the fusion experience O_{green} does not. If one is seeking to learn of the book's colour, it is reasonable to privilege the former over the latter. Privilege should also be granted to experiences in which *only* the book is contributing to the colour experience along a LOS.[9] Call experiences in which only one object is contributing to the colour experience along a LOS, $LOS_{1\text{-}object}$ experiences. Thus, to use fusion cases to ground Relationalism one must exclude these alternatives, for with these alternatives in play, *No Privilege* and then *Ecumenicism* fail.

We should therefore resist the reasoning for *Explanatory Import* in LOS scenarios in which layering *or* fusion occurs. If experienced layering occurs, then the kind of perceptual relativity at play is not the kind needed to sustain *Explanatory Import*. If one's aim is to learn about an object's colour and experiential fusion occurs involving that object and another object along the same LOS, such fusion experiences are epistemically inferior to $LOS_{1\text{-}object}$ or layering experiences involving that object.[10]

---

[9] For example perhaps along a LOS is the book, which is illuminated by an illuminant I_1, and viewed through some air V_1. Consider a context in which *only* the book's contribution to colour is experienced along that LOS. If one seeks to learn of the book's colour, this context should be privileged over a context that fuses together contributions from the book, illuminant and/or air.

[10] Note that even with fusion experiences there may be a means of "inferring" an object's non-relational colour, namely, by trying (intellectually) to isolate the various contributions to experienced colour that an object makes across a host of fusion experiences.

In addition a host of intermediary cases between "perfect" experienced layering and "perfect" experienced fusion are possible. There might be cases close to but note quite achieving perfect layering, where the opaque and transparent object colours are almost perfectly distinct in experience but a perceptible degree of mixing occurs. On the other extreme there might be cases close to but note quite achieving perfect fusion, where the opaque and transparent object contributions to experienced colour are almost perfectly fused, but a residual layering effect remains. And so on (see xxx for some additional options). Such variations do not challenge the above analysis. To my mind the only means of undermining the above analysis is via a demonstration that experiences of colour layering and of $LOS_{1\text{-}object}$ are impossible or bound toward being so. I know of no attempt to provide this demonstration, but welcome any such efforts.

§3.2 Layering: its potential reach and confounds.

Given the sheer size of the class of LOS relativities (see the beginning of §3), the above argument marks a potentially wide-reaching means of undermining the *Uniformity Hypothesis*. Its potency hinges on the proliferation of non-fusion ($LOS_{1\text{-}object}$ or layering) colour experiences within the LOS class. If this is the case, why don't we start scrutinizing? Unfortunately, assessing instances of LOS relativities for layering/$LOS_{1\text{-}object}$/fusion yields several *confounds*. Here, in brief, are four of them. For simplicity I focus on layering and fusion experiences, and remark on $LOS_{1\text{-}object}$ experiences at the end of the section. It will help to roughly categorize the many LOS relativities into three *classes*, colour perceptions involving contributions from transparent *films*, *volumes*, and *illuminants*.

Heterogenous classes. We should be wary of any attempts to conclude that a particular class contains experiences of layering to the exclusion of fusion, or vice versa. There is nothing to prevent one context involving a transparent film from yielding an experience of colour layering, and another from yielding experiential fusion. We know that our perceptual systems will have much to say about whether experiential layering or fusion occurs (see, e.g., references in note 3). Stated most generally, there are various factors within a complex scene that prompt, to some extent or other, one's visual system into a state involving experiential fusion or layering. Nothing as simple as "perceptions involving transparent films yield layered

experiences" is antecedently credible. We instead must be open to the possibility that some members of each perceptual class contain experiential fusion and some others layering.

Perceptual ambiguity. A stimulus in one of these classes might give rise to a perceptual ambiguity that "flips" between experiential fusion or layering. For example a partly shadowed wall might at one moment be (correctly) experienced as differently illuminated and similarly painted, and at another moment (incorrectly) as similarly illuminated but differently painted. In the first experience the illuminant and surface contributions to colour are experientially distinct, and in the second at least some of the illuminant and surface contributions are fused together. Arguably experiential layering occurs in the first instance in ways that are replaced with experiential fusion in the second.[11]

There can also be ambiguities across classes, say in which a volume perception (e.g., looking at a scene through thick smog) is experienced as a member of that class then as a member of the illumination class (e.g., as a scene under low levels illumination). In principle there is nothing to prevent a single stimulus from being involved in both types of ambiguity.

First-personal limitations. There are cases in which it is fairly obvious that experiential layering versus fusion is occurring. However, the boundaries of these kinds of experiences is undoubtedly vague, yielding a range of cases with only subtle experiential differences. Further, subjects are in general not accustomed to making judgements about layering/fusion. This is a troublesome combination. Imagine asking a subject: "Are all colour elements experienced along this LOS attributed to that opaque thing (fusion) or are some colour elements instead attributed to a transparent intermediary (layering)?" We cannot antecedently expect straightforward data from subjects presented with a range of subtly difference cases. To take a concrete example, when one's experience when looking at the world gets a yellowish tinge due to the yellowing of the lenses of one's eyes with age, is the experienced yellowing an instance of attributing the yellowness to objects in the world (fusion), or of looking at those objects through a yellow lens (layering)? If the former, then Relationalism is potentially critical to explaining the phenomenon, but if the latter then it is otiose. A subject's

[11] There is much more to be discussed about this kind of case, but space demands a brief treatment.

merely indicating some form of "things look more yellow" does not decide the matter. Our access to the fine-grained nature of our own colour experience is not adequately reliable to straightforwardly assess the matter.

Colour constancy. The potential impact of colour constancy on this discussion is immense. When colour layering occurs, and one of the layers varies while the other remains constant, and the subject experiences just that, then a form of colour constancy obtains. For ease of reference call it *layered constancy*. As noted in §1 (and discussed at length in xxx) this model for constancy has potential instances in constancy cases involving opaque surfaces and variations in filters, or volumes, or illuminants (and converse cases where the opaque surfaces vary and the other elements remain stable). Given that constancy obtains throughout our daily lives, the potential number of instances of layered constancy is vast. Each such instance reduces the extent to which Relationalism should be used as an explanatory tool.

Cohen (2008) details a Relationalist-friendly account of colour constancy aptly deemed a *counterfactualist* account. According to it, the contributions that the constant and variable elements make to colour experience are fused together to yield a relational colour that is both experienced by the subject and in the world. Thus, during constancy cases there is in fact no constant colour, only colour variation. Take for example a partly shadowed blue wall. The surface and illuminant contributions to experience are fused together on both the shadowed and unshadowed parts thus yielding two different colours, a lighter and darker shade of blue. While there is blue on both parts of the wall, and thus a constancy of colour category, there is no constant fine-grained colour. Since these distinct fine-grained colours of the wall are defined by reference to both surfaces and illuminants, they are relational in at least this regard.

That is the "occurrent" (to use Cohen's term) experiential and ontological reality of constancy cases. To make this an account of colour constancy, however, something has to be added to explain those aspects of subjects' behaviours that indicate that there is more to these cases than merely colour variations, that there is a sense in which the relevant colour samples are constant. For this purpose Cohen postulates the existence of *counterfactual* colour contents, contents to the effect that the colours in these two parts of the world would be the same were the perceptual conditions the same. For example the blues on these two parts of the wall would be the same were illumination conditions the same. The thought is that our visual-cognitive systems

are making these counterfactual judgements whenever colour constancy obtains (which is throughout our lives), and our commitment to constant colours in these cases is an expression of these contents, not an expression either of a constant experienced colour or of a constant instantiated colour. I will make two remarks.

The first remark is methodological. A counterfactualist analysis of constancy is not empty, on the contrary I suspect it has instances. However, one's goal should not be to try to force all instances of colour constancy into one model (be it a layering, counterfactualist or some other model) to accommodate one's broader objectives (be them a defense of colour objectivism or Relationalism). Our goal should be to examine various instances of colour constancy with various models in mind to try to uncover which one or more of them best explain individual or classes of cases. The potential result is that colour constancy is a varied phenomenon, having some instances of layered constancy, others of counterfactual constancy, and so on.

My second remark is the now-familiar one that, with regard to colour constancy, Relationalism only has explanatory import where instances of counterfactualism obtain, and has no explanatory import otherwise. Since colour constancy is such a widespread phenomenon, occurring throughout scenarios involving LOS relativities, there is no simple, straightforward way for Cohen to argue that Relationalism explains these cases, and likely no way in general, since layered constancy plausibly has instances.

Here is the general lesson: there are some and potentially numerous instances of experienced colour layering; determining their true extent must be done not by asserting the *Uniformity Hypothesis* or any such theoretical principle, but by careful examination of relevant cases; due to various confounding factors that examination will not be straightforward. In other words, we cannot assert the *Uniformity Hypothesis*.

Note that this conclusion follows regardless of the prevalence of $LOS_{1\text{-}object}$ experiences. Recall that an $LOS_{1\text{-}object}$ experience includes a case in which a surface is illuminated by an illuminant, and viewed through some air, but in which *only* the surface contribution to colour is experienced along that LOS. In one sense the independence of the above conclusion from the prevalence of $LOS_{1\text{-}object}$ experiences is fortunate, because I am doubtful that $LOS_{1\text{-}object}$ experiences are prevalent. I instead suspect that illuminants and transparent

volumes make regular contributions to the colours experienced along LOSs. However, if it turns out that $LOS_{1\text{-}object}$ experiences are prevalent, then our confidence in the *Uniformity Hypothesis* is further eroded.

§4 Conclusion

The most compelling argument for Relationalism infers the inherently relational character of colour from cases in which one's experience of a colour contextually depends on one's experience of other colours. I have argued that this reasoning has limits, that in cases involving experienced colour layering, the contextual interdependency of one's experience of one colour on one's experience of another does not ground an inference to the relational nature of colour. However, I have *not* argued that this argument for Relationalism is generally bankrupt. The prevalence of a rich phenomenon like center-surround contrast effects should tempt us all to be Relationalists of some sort. These two conclusions are a way of expressing what results from denying the *Uniformity Hypothesis*. The relativities of colour, rather than yielding "identical considerations" that put "pressure on us to favor a uniform reaction in each case", yield differing considerations, some of which pressure us to become Relationalists, others of which do not (Cohen, 2009, 20-21). The global pressure toward Relationalism is illusory, and to the extent that we succumb to it, we risk being blind to surprising possibilities, like the potential ubiquity of colour layering.

§5 References

Anderson, B.L. (2008). Transparency and Occlusion. In Basbaum, A.I., Kaneko, A., Shepherd, and Westheimer, G., eds., *The Senses: A comprehensive Reference, Vol. 2, Vision II*, Albright, T.D., & Masland, R. (volume eds.), pp. 239-244. San Diego, USA: Academic Press.

Brainard, D., W. Brunt & M. Speigle (1997). Color constancy in the nearly natural image. I. Asymmetric matches. *J. Opt. Soc. Am., 14*(9): 2091-2110.

Brainard, D. & Wandell, B. (1988). Classification measurement of color appearance. *Investigative Ophthalmology & Visual Science, 29*: 162.

Byrne, A. & D. Hilbert (2003). Colour realism and colour science. *Behavioral and Brain Sciences, 26*: 3-64.

Casati, R. (2009). Are shadows transparent? An investigation on white, shadows and transparency in pictures. *Res,* 56: 329-335. My source is the Author Manuscript ijn_00171296, pp. 1-12, Sept 2007.

Cohen, J. (2008). Colour constancy as counterfactual. *Australasian Journal of Philosophy,* 86: 61-92.

-- (2009). *The Red and the Real: An Essay on Colour Ontology.* Oxford, UK: Oxford University Press.

Jameson, D., & Hurvich, L. (1989). Essay concerning color constancy. *Annual Review of Psychology,* 40: 1-22. Reprinted in Byrne & Hilbert, eds., 1997, *Readings in Colour, Vol. 2: The Sciene of Color*, pp. 177-98. Cambridge Mass: MIT Press.

Khang, B.-G. and Q. Zaidi (2002a). Accuracy of color scission for spectral transparencies. *Journal of Vision* 2(6): 451-466S.

-- (2002b). Cues and strategies for color constancy: perceptual scission, image junctions and transformational color matching. *Vision Research* 42(2): 211-26.

Kingdom, F. A. A. (2011). Lightness, Brightness and Transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research*, 51: 652-673.

Kitaoka, A. (2005). "A new explanation of perceptual transparency connecting the X-junction contrast-polarity model with the luminance-based arithmetic model". *Japanese Psychological Research*, 47: 175-187.

Land, E. (1986). Recent advances in Retinex theory. *Vision Research, 26*: 7-21. Reprinted in Byrne & Hilbert, eds., 1997, *Readings in Color, vol. 2: The Science of Color*, pp. 143-160. MIT Press: Cambridge, Mass.

Land, E. & McCann, J. (1971). Lightness and Retinex theory. *Journal of the Optical Society of America, 61*: 1-11.

Metelli, F. (1974). The perception of transparency. *Scientific American*, 230(4): 91-98.

Shevell, S. K. & F. A. A. Kingdom (2008). Color in Complex Scenes. *The Annual Review of Psychology 59*: 143-166.

Wandell, B. (1989). Color constancy and the natural image. *Physica Scripta, 39*: 187-92. Reprinted in Byrne & Hilbert, eds., 1997, *Readings in Color, vol. 2: The Science of Color*, pp. 161-75. MIT Press: Cambridge, Mass.

Wollschläger and Anderson (2009). The role of layered scene representations in color appearance. *Current Biology, 19*: 430-435.