COLOUR RELATIONS IN FORM

Will Davies

St Anne's College, University of Oxford

Abstract: The orthodox *monadic determination thesis* holds that we represent colour relations by virtue of representing colours. Against this orthodoxy, I argue that it is possible to represent colour relations without representing any colours. I present a model of iconic perceptual content that allows for such primitive relational colour representation, and provide four empirical arguments in its support. I close by surveying alternative views of the relationship between monadic and relational colour representation.

Suppose you view a banana next to a strawberry. You perceptually represent the banana as yellow and the strawberry as red. You also perceptually represent the banana as yellower than the strawberry, and the strawberry as redder than the banana. The orthodox view is that you represent these relations by virtue of representing the monadic colours of the banana and the strawberry. This view is intuitive, and deeply rooted within the philosophy of colour (Section 1). Against the orthodoxy, I argue that it is possible to represent colour relations without representing any colours.[1] Section 2 presents a model of iconic perceptual content that allows for such primitive relational colour representation. Section 3 provides four empirical

---

[1] 'Colour' here means *chromatic colour*.

arguments in its support. Section 4 surveys alternative views of the relationship between monadic and relational colour representation.

A brief clarification, before we begin. My argument concerns our visual perceptual representation of colour relations, rather than visual experience per se. I explain my reasons at length elsewhere (Davies, 2018). By 'visual perceptual representation,' I mean a state of the visual perceptual system that is type-individuated by its accuracy conditions. An example of a relational colour representation, then, would be a state type-individuated by the accuracy condition that the banana is yellower than the strawberry. I assume, minimally, that some aspects of the character of our visual experience of colour relations supervene on such representational states. As such, some of my arguments appeal to differences in visual phenomenology to support claims about relational colour representation. Other arguments, however, exclusively appeal to neural and functional facts about the colour vision system.

1. The Monadic Determination Thesis

Evie wants to paint a picture of a banana next to a strawberry. She especially wants to depict that the banana is yellower than the strawberry, and the strawberry redder than the banana. How will Evie represent these colour relations? Easy! She applies yellow paint to represent the banana here, red paint to represent the strawberry there, and ipso facto represents that the banana is yellower than the strawberry. Indeed, this seems like the *only* way for Evie to represent such relations. Painting is a monadic business: her palette contains red and yellow paints, but no 'redder than' or 'yellower than' paints.

The orthodox view is that Evie's visual perceptual representation of colour relations works in a similar way. She perceptually represents the banana as yellower than the strawberry, and the strawberry as redder than the banana, by virtue of perceptually representing the banana as yellow and the strawberry as red. More generally, our perceptual representation of monadic

colours determines our perceptual representation of colour relations. Call this the *Monadic Determination Thesis* (MDT).

The MDT is an intuitive view. By analogy, it is natural to think that perceptually representing objects as having particular spatial locations determines our representation of their spatial relations. Equally intuitively, perceptually representing the crash and the bang as occurring at times t1 and t2 determines our representation of the bang succeeding the crash. Part of the MDT's appeal derives from the well-worn, putative, analogy between visual perception and painting. Locke (1975: II, ix, 8) remarked that the 'ideas we receive by sensation' are of a 'plane variously coloured, as is evident in painting.' Hume (2000: 56) claimed that 'all bodies, which discover themselves to the eye, appear as if painted on a plain surface.' Just as painters represent objects via dabs of coloured paint on a surface, for Locke and Hume, visual representations comprise an array of 'sensible points,' or 'atoms or corpuscles' endowed with colour. If visual perception is structured this way, the MDT seems inevitable.

The MDT also coheres with the standard metaphysical view that objects' intrinsic properties determine their relational properties. As Ladyman (2014: §4) notes,

The idea that there could be relations which do not supervene on the non-relational properties of their relata runs counter to a deeply entrenched way of thinking… [I]t is often assumed that a structure is fundamentally composed of individuals and their intrinsic properties, on which all relational structure supervenes.

In the case of colour, moreover, Pautz (2006: 554) argues that

It seems perceptually obvious that the fact that blue resembles purple more than green obtains by virtue of the intrinsic character of the colours. To see that the colours fall into this resemblance order, you only need to focus on the colours themselves.

The standard metaphysical view is 'perceptually obvious', Pautz claims, because we can see that blue resembles purple more than green, just by virtue of seeing blue, purple, and green.[2] As should be clear, Pautz's justification for the ontological thesis assumes the MDT.

There are two main ways to develop the MDT. Developed metaphysically, the view is that monadic colour representations fully ground relational colour representations.[3] The view implies that, necessarily, any change in relational colour representation entails a change in monadic colour representation. Developed causally, the view is that our cognitive architecture contains mechanisms taking monadic colour representations as inputs, yielding relational representations as outputs. As a toy example, suppose that Evie's visual system represents the strawberry as red$_{23}$ and the banana as yellow$_{17}$. For convenience, assume that perceptually representing a monadic colour involves attributing proportions or magnitudes of the elemental colours red, green, yellow, and blue.[4] Our hypothetical mechanism 'reads off' the proportions of redness and yellowness attributed to the strawberry and banana, computes their comparative values, and outputs the representation that the strawberry is redder than the banana.

The model of perceptual content developed below undermines both readings of the MDT. The model shows how a perceptual state might represent colour relations without

---

[2] I discuss issues regarding colour similarity and colour ontology in Davies (2014).
[3] Papineau (2015) calls this 'the orthodox view'; Morrison (2015) labels it 'colour atomism'. I discuss connections with these authors in Section 4.
[4] See Byrne & Hilbert (2003) on hue magnitudes .

representing any colours. Against the metaphysical MDT, such cases establish the possibility of changes in relational colour representation without changes in monadic colour representation, because no monadic colour representation. The case against the causal MDT is more complex, and warrants independent development. Roughly, visual systems to which this model applies, in some circumstances, could compute relational colour representations without having computed any monadic colour representations.

2. The Chromatic Edge Model

This Section introduces a model of perceptual content that allows for the primitive representation of colour relations: the *chromatic edge model*. Regarding product differentiation, Mandik (2014: 230) argues that linguaform models of perceptual content may allow cases where subjects perceive an *F*-relation between *x* and *y*, without attributing any particular *F*-property to *x* or *y*. My view is novel in that it develops a model of iconic, rather than discursive, representational content. This is a harder task, as theories of iconic representation typically follow the empiricists in assuming that the syntactic primitives are 'points' or 'pixels' with monadic colour properties.[5]

2.1 Topographic Maps

Advocates of iconic theories of perceptual content sometimes motivate the view by noting that visual representation is more like that of maps than natural language.[6] I develop this analogy by considering topographic maps, which represent terrain or relief using *contour lines*, curves connecting contiguous points of equal elevation. The map legend will include a *contour interval* specifying the difference in elevation between each successive contour line.

---

[5] Kosslyn et al. (2006: 11-12): 'symbols belong to two form classes: points and empty space… The points can vary in size, intensity, and colour.'

[6] Burge (2010: 540): 'the elements in perception are organized nonpropositionally. Let us suppose that they are organised in ways that are structurally isomorphic with a topological or geometrical structure. Think of a map… In vision, the elements in perception have something like the form of a map or sketch from an egocentric perspective.'

Standard contour lines indicate increases in elevation; hachured contour lines indicate decreases in elevation. In Figure 1, contour lines A, B, C, and D join points of equal elevation, with a contour interval of 10m.[7] The points joined by B therefore are 10m higher than the points joined by A. The points joined by the hachured line D are 10m lower than the points joined by C.
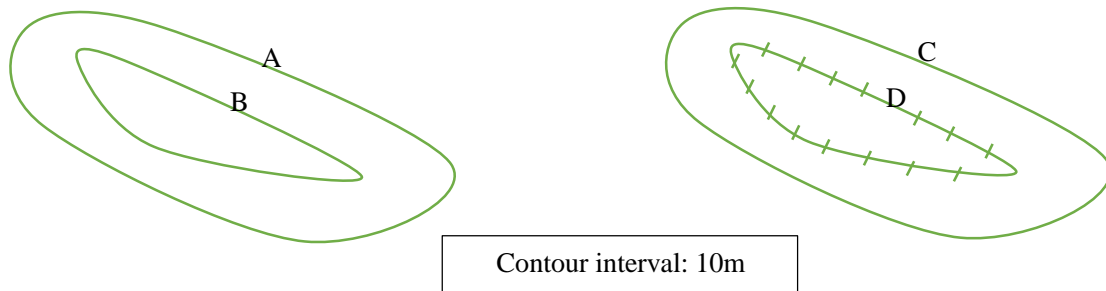

Contour interval: 10m

Figure 1

Contour lines represent comparative elevations between two or more sets of points, without representing the particular elevations of those points. In the diagram above, the contour lines suffice to determine the comparative heights of the points connected by lines A and B, and lines C and D. In practice, topographic maps normally include *reference* contour lines labelled with specific elevations. Labelling A as 130m, for example, would enable us to calculate that B joins points at 140m. This is crucial information in navigation. Such references are not mandatory, however, and contour lines per se do not represent monadic elevation properties.

Now consider a simple colour topography map, as in Figure 2b. The map represents the scene in Figure 2a via colour contour lines that enclose regions with coordinates corresponding to the spatial locations of the disc and square. The colour contours also represent the colour relations between regions. Schematically, a contour with colour *C*, which segregates a region from its immediate surround, represents that the region is *C-er* than its

---

[7] Please see the online version of this paper for colour images.

surround. In Figure 2b, the red contour represents that the disc is redder than its immediate

surrounds, and the yellow contour represents that the square is yellower than its immediate

surrounds. We could complicate the syntax by allowing variations in the width of colour

contours, for example, to represent variations in the magnitude of colour variation, with

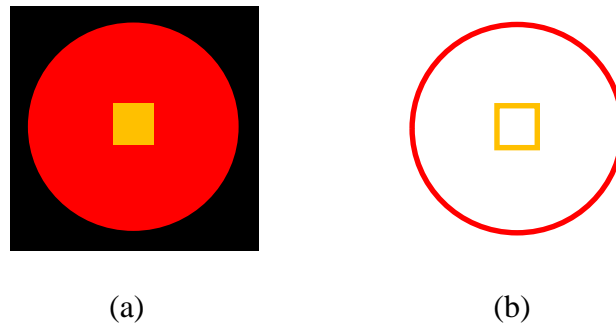thicker lines corresponding to larger magnitudes.

(a)   (b)

Figure 2

2.2. Edge Representation in Visual Perception

Our simple colour topography map represents directions and even magnitudes of colour

contrast via features of colour contour lines. The edges that we perceive obviously are not like

this: they lack width, and do not appear to have monadic colour properties. Nevertheless, edges

do visually appear in certain ways to us. With Casati and Varzi (1999: 71), I assume that edges

and other boundaries are '*bona fide* spatial entities' that 'enter the content of our perceptions.'

I focus on edge representations with iconic formats. Iconic representations are *holistic*,

in that they lack canonical decompositions into syntactic constituents that separately denote

individuals and features.[8] Consider a photograph of Big Ben. There is no way to decompose

this photograph into constituent elements that separately encode the individual, Big Ben, and

its features, such as its height, colour, and texture. Any constituent part of the image represents

a part of Big Ben, but also represents the features of that part. Contrast this with a discursive

---

[8] Quilty-Dunn (2017: 63).

representation such as 'Big Ben is tall,' which has a canonical decomposition into the primitives 'Big Ben' and 'tall', separately denoting an individual and a feature.

I assume, then, that an iconic visual representation of an edge will holistically encode features such as – at a minimum – its curvature, length, and orientation (C.f. Burge, 2014: 493). In addition, I assume that an edge between two regions will be represented as belonging to, or 'owned' by, just one of those regions. In this respect, edge representation has a classical topology. In viewing Figure 2a, for example, we perceptually represent the edge between the disc and the black background as belonging to the disc, and the edge between the square and the disc as belonging to the square. These representations are part of the standard explanation of the figure-ground organisation of the scene, wherein the disc appears as figure against a black ground, while the square appears as figure against a red ground.[9] One might object that we sometimes perceive boundaries between objects lacking an obvious figure-ground organisation. Does the edge between the squares in Figure 3 appear to belong to the red or yellow square? The answer is unclear.[10] Perhaps we will have to abandon the classical topology in such cases. Here I proceed on the simplifying assumption.
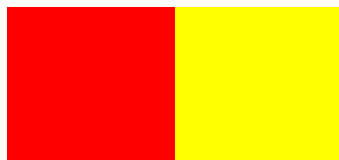

Figure 3

2.3 Directed Chromatic Contrast Properties

On my view, visual edge representation attributes chromatic features, along with geometric features such as curvature and length, and topological features such as ownership relations. These features are *directed chromatic contrast properties*, or *chromatic contrasts* for short.

---

[9] I discuss the role of edge representations in figure-ground in Davies (m.s.).
[10] Smith (1998/99), Casati (2003).

The chromatic contrast of an edge specifies the type of chromatic variation obtaining between the regions bounded by that edge. Take the edge between the yellow square and the red disc in Figure 2a. We can gloss its chromatic contrast as: *owned by a region that is n units yellower than its immediate neighbouring region.*

We can use *chromatic contrast vectors* to model the content involved in representing such contrast properties. To simplify, assume a standard three-dimensional opponent colour space, with x, y, and z axes corresponding to red-green (RG), yellow-blue (YB), and black-white (BW) dimensions.[11] Neutral grey is mapped to the point $(0, 0, 0)$. We represent increases in redness, yellowness, and blackness via increases in x, y, and z values respectively, and represent increases in greenness, blueness, and whiteness via decreases in x, y, and z values respectively. We can then represent colour variations via vectors of the form $\begin{pmatrix} a \\ b \\ c \end{pmatrix}$, where positive a, b, and c, values represent magnitudes of change in the direction of redness, yellowness, and blackness respectively, and negative a, b, and c, represent magnitudes of change in the direction of greenness, blueness, and whiteness respectively. Thus for example $\begin{pmatrix} 8 \\ 0 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 17 \\ 0 \\ 0 \end{pmatrix}$ represent different magnitudes of variation in the direction of redness; $\begin{pmatrix} 8 \\ 0 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} -8 \\ 0 \\ 0 \end{pmatrix}$ represent variations of equal magnitude in the opposing directions of redness and greenness; $\begin{pmatrix} 8 \\ 0 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 8 \\ 0 \end{pmatrix}$ represent variations of equal magnitude in the orthogonal directions of redness and yellowness; and so on.

Think of these vectors as mappings from one location (though no particular location) in colour space to another. As noted above, edges are represented as belonging to one, and only one, of the regions bounded by the edge. When a contrast vector is attributed to an edge,

---

[11] I do not ultimately think that the familiar opponent colour space is appropriate to model the chromatic contrasts attributed in visual edge representations. I use the opponent colour space here and throughout for heuristic purposes.

the syntax of this representation dictates that the vector always 'points' in the direction of the region that owns the boundary. To illustrate, suppose that we attribute $\begin{pmatrix} 17 \\ 0 \\ 0 \end{pmatrix}$ to the circular boundary in Figure 2a. We thereby represent that the disc is 17 (arbitrary) units redder than its background, because we represent this boundary as being owned by the disc, rather than the background.

Regarding precedents, Clark (2000: 126ff) distinguishes 'occupant coding' of colour, which represents monadic point values, and 'difference coding', which codes chromatic contrasts. Clark's (2000: 130-1) view is that 'the informative or content-bearing elements are not the swaths of constant hue, but the edges, where things change.' He suggests that 'detecting colour differences can help identify what type of edge it is: as one where spectral properties change... [T]he suggested content of chromatic sensory states is something like "here things get redder,"' (2000: 128, fn. 38). My view cashes out this suggestion in terms of the attribution of chromatic contrast vectors to edges, within a topologically structured representation that encodes boundary-ownership relations. The following Section also presents four new arguments in its support.

3. Arguments for the chromatic edge model

3.1. Liebmann Effects

Figure 4

In colour contrast cases, the appearance of an object's monadic colour varies with changes in surrounding context. In Figure 4, the central rectangles in each pair of images have the same surface spectral reflectance properties, but appear strikingly different depending on the colour of the surround.[12] Less frequently noted, but equally obvious, is that differences in chromatic contrast also influence the appearance of edges. Consider the two purplish rectangles against the blue and green backgrounds. The edges of the left rectangle appear less distinct than the right rectangle. The same applies to the right greenish rectangle against the cyan background, whose edges appear less pronounced than the left greenish rectangle against the yellow background. Moreover, to my eye, the blue-green edge in the top left of Figure 4 differs in quality from the red-blue edge in the top right. Finally, Figure 5 shows two series of decreasing contrasts between the central squares and surrounding discs. While the edges of the squares are clearly visible in the leftmost images, they become increasingly indistinct as we progress right, finally becoming invisible.

---

[12] Such effects provide important motivation for the relationalist colour ontology of Cohen (2009).
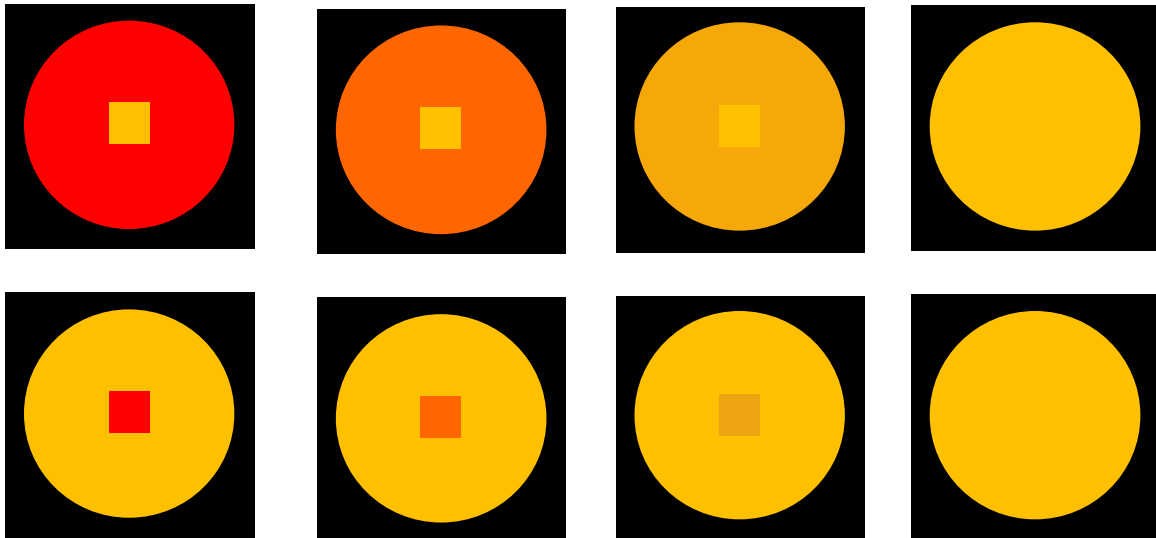
Figure 5

From these observations, it seems that both *magnitudes* and *directions* (in colour space) of chromatic contrast may influence the visual appearance of edges. Gestalt psychologist Susanne Liebmann (1927, reproduced in West et al. 1996: 1461) examined these influences by assessing the appearance of borders between equiluminant stimuli,

How do figures appear when they are different in colour but equal in brightness to the ground? In which way do colours without a brightness difference contribute to the perceived existence of figures, bodies, things, and the structure of the visual environment?

She describes the now eponymous *Liebmann Effect* (1996: 1461),

When … figures are perceived, then generally everything is soft, jellylike, colloidal. The figure does not stand out from the ground. It is not only there with blurred

contours, but the whole of it is diffuse and ill-defined. Drastic figural changes occur…

There are in psychology only a few phenomenal changes of so drastic a nature.

Some authors have taken these effects to indicate that colour vision cannot support any sort of form segmentation.[13] As Mollon (1989: 30) notes, however, this overplays the extent of the effects, which vary with the type of colour contrast – along with other factors.[14] In a minority of cases, figural boundaries are lost altogether. Otherwise, subjects perceive edges of varying quality and distinctness. In approximate order, contours are 'softest' for equiluminant blue-green pairs, followed by blue-yellow, green-yellow, yellow-red, red-green, and blue-red (1996: 1474-6).[15] In general, the effects are stronger for lower saturation colours. Note that edge softness/hardness is not a simple function of the 'proximity' or contrast magnitude between the two colours: the direction matters. For example, equally discriminable blue-green and red-blue pairs form borders that appear different in softness/hardness.

The chromatic edge model provides a good explanation of these effects. The apparent 'softness' or 'hardness' of an edge seems to vary with the magnitude and direction of chromatic contrast between the regions bounded by that edge. As noted above, I assume that edges are represented in visual perception, and that discriminable differences in edge appearance license inferences about differences in edge representation. The noted differences in edge appearance thus indicate differences in edge representation that vary with magnitudes

---

[13] Livingstone & Hubel (1987: 3444; 1988: 747).

[14] For example, the effects are stronger for small, narrow, figures than large, broad figures, and stronger for complicated than simple figures.

[15] C.f. Koffka's (1936: 127) distinction between 'hard' colours such as red, which 'segregates best,' and 'soft' colours such as blue, which 'segregates least.' Interestingly, opponent colour pairs of blue-yellow and red-green do not form the 'hardest' edges. Red-green edges are among the hardest, but no harder than non-opponent pairings of blue-red. Blue-yellow edges are among the softest. I reemphasise that the familiar opponent colour space is thus unlikely to be appropriate for modelling chromatic edge content.

and directions of chromatic contrast. It is plausible, then, that these differences in edge representation are, straightforwardly, differences in the representation of magnitudes and directions of chromatic contrast.

One might object that this is not how edges look to us. Edges look smooth, jagged, soft, hard, distinct, indistinct; they do not look *red-blue-contrast-type* or *red-green-contrast-type*. In response, though, who says that edges do not look such ways to us? Clearly, we don't ordinarily use chromatic contrast terms to describe edges. But this directly undermines my view only if we assume that all properties attributed in visual perception must be conceptualised by the subject. I reject that assumption. Moreover, in light of the Liebmann experiments, it seems not implausible that a subject could *come* to describe edges in such ways, given suitable motivation. The following case presses this point.

Suppose we train a subject to become expert in discriminating differences in edge qualities between pairs of equiluminant colour stimuli. Their samples include figures of various shades of blue against backgrounds of various shades of green, and vice versa; red figures against yellow backgrounds, and vice versa; and so on. Once they reach ceiling on this discrimination task, we train them on an identification task. For every discriminable type of edge quality, they learn which (ordered) pairs of colours form edges of this type. Given a $blue_{14}$ figure on red_{31} background, for example, they can identify the edge as having the same quality as $blue_{13}$ on red_{32}, but also (say) $purple_5$ on $orange_2$, and so on. We then encourage them to form cognitive categories for these different quality types, and introduce names for them. They may then say, of a given edge, 'that looks like a *such-and-such* edge,' where 'such-and-such' denotes a category that includes $blue_{13}$ on red_{32}, $purple_5$ on $orange_2$, and so on. These categories might not map neatly onto familiar categories such as *red-green-*

contrast: the categories would likely be sui generis.[16] Nonetheless, this subject would describe how edges look using categories of chromatic contrast. If you find this believable, you should be open to the idea that edges have distinctive contrast-dependent looks.

3.2. Cerebral Achromatopsia

Cerebral achromatopsia is a selective impairment in the ability to see monadic colour properties, resulting from cortical rather than retinal abnormalities. Complete cerebral achromatopsics seem to have entirely achromatic visual perception. They fail standard diagnostic tests for colour vision, involving discrimination of isoluminant stimuli, colour naming, and ordering samples of colour chips. Despite these failures, studies reveal a surprising amount of spared processing within the colour vision system. A complete achromatopsic known as 'MS' was found to have a spectral sensitivity function consistent with preserved colour-opponent processing. He could detect contours in sinusoidal gratings made up from isoluminant hues, coloured shapes embedded in static or dynamic grey backgrounds, and could detect the direction of movement of an isoluminant chromatic grating (Cowey & Heywood, 1995).

Several commentators have suggested that cerebral achromatopsics perform these tasks by virtue of preserved 'colour-for-form' processing.[17] A first indication was that some achromatopsics could detect the figures in Ishihara colour plates at a distance of two metres, though not at normal reading distance. At normal reading distance, the dominant contrasts in these plates are the varying luminance contrasts between the small coloured discs. As Mollon and colleagues (1980: 133) note, however, at two metres 'the luminance contours of individual discs are no longer resolved and the dominant contour is the hue boundary between figure and ground.' Patients report seeing a figure embedded within the display and

---

[16] See footnotes 11 and 15.
[17] Shevell & Kingdom (2008: 152), Chirimuuta & Kingdom (2015: 226).

can accurately name it and trace its outline, though they report not seeing the colours of the discs.

Subsequent studies have further probed this spared capacity. Kentridge and colleagues (2004a: 822) sought to assess 'the nature of the local chromatic contrast signals that can be accessed and discriminated by the visual system independently of their role in the perception of constant surface colour.' They strikingly claim that 'a cortically colour blind observer [MS] *perceives* chromatic local-contrast signals,' (2004a: 829, italics added). One discrimination task sought to test whether MS could discriminate magnitudes and directions of chromatic contrast. MS had to identify the odd one out among three coloured discs, presented against a uniformly coloured background, as in Figure 6. For example, image (b) consists of three green discs against a yellow background. Two of the discs were identical, with cone-contrasts of 27.5% from the background. The third disc varied between cone-contrast increments of 75%, 62.5%, or 50% from the reference contrast. While clearly abnormal, MS was able to distinguish the odd one out at all three contrast increments, though his performance was worse for the smaller increments (2004a: 825-6). They also found that MS could discriminate contrasts of equal magnitude in the opposing directions of redness and greenness, as in image (d). MS viewed three discs with equal cone-contrasts of 27.5% from the yellow background, but two of the discs were red (green) and one of the discs was green (red). MS successfully discriminated the odd disc on 80 out of 96 trials, considerably better than chance (2004a: 826).
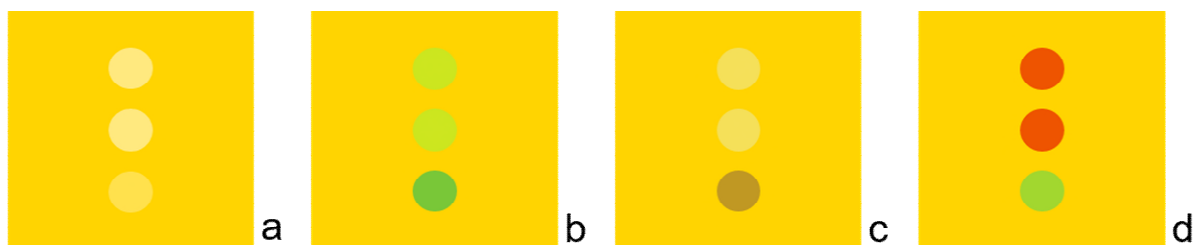

Figure 6 (Reprinted with permission from Kentridge et al. 2004a)

Tellingly, Kentridge and colleagues attribute MS's task performance to his ability to discriminate differences in the appearance of the edges of the discs. MS himself reports that the edges of the discs look different, and that this is how he can tell the odd one out. This fits a wider pattern. In an earlier study, MS discriminated sequences of isoluminant squares ordered in respect of colour, from sequences of randomly ordered squares. Heywood and colleagues (1991: 802) reported that MS's 'verbal replies showed that he did so by detecting an edge between two stimuli that were, to him, perceptually identical.' They claimed that these edges differed in 'salience' for MS, though they did not explain this notion.

In further support of the edge-based interpretation, MS's performance on the task fell to chance when the discs were surrounded by black annuli, as in Figure 7. This issue requires considerable care. A first gloss is that the annulus obscures the edge between disc and background, concealing the chromatic contrast information underlying MS's discriminations. Note, however, that the image still contains differing chromatic contrasts between the red and green discs and the black annuli. Why cannot MS discriminate these contrasts? Kentridge and colleagues (2004a: 828) surmise that MS's 'access to these [chromatic edge] cues is…. disrupted by coincident luminance contours.' That is, whereas MS can discriminate the 'pure' chromatic contrasts between equiluminant discs and background, he seems unable to discriminate 'mixed' chromatic and luminance contrasts between discs and annuli, and annuli and background.

Figure 7 (Reprinted with permission from Kentridge et al. 2004a)

Now, although this evidence supports an edge-based explanation of MS's abilities, it actually threatens to undermine the chromatic edge model. The model proposes that edge representations have chromatic contrast content. At first pass, however, MS's inability to discriminate the chromatic contrasts in Figure 7 suggests that his edge representations *conflate* chromatic with luminance contrast. Indeed, Kentridge and colleagues (2004a: 828) hypothesise that MS's edge representation is related to the activity of cells with 'mixed colour-luminance opponency', which are 'good candidates for the mediation of the perception of form from colour.' They reason that,

> [I]f the luminance sensitivities of the centre and surround balance then the cell will respond exclusively to the chromatic contrast of a border. If, however, the luminance sensitivities of centre and surround differ, then the cell will respond to both chromatic and luminance borders.

Placing this in broader context, a common criterion for the possession of colour vision is the ability to discriminate differences in light wavelength, independent of changes in light intensity.[18] If these variables are confounded in the subject's visual response, then they do not possess colour vision. It seems, however, that MS is unable to discriminate differences in chromatic contrast independent of changes in luminance contrast. By parity of reasoning, we may doubt whether MS can visually represent chromatic contrast.

In response, it is important not to overstate the extent to which MS's visual representations conflate chromatic and luminance contrast. It does seem that MS cannot

---

[18] For detailed discussion of such discrimination-based criteria for colour vision, see Davies (2018).

Forthcoming in *Philosophy and Phenomenological Research*. Please cite published version.

14/01/2020

discriminate chromatic contrasts coincident with very large luminance contrasts, as with the black annuli. MS is nonetheless able to discriminate chromatic contrasts coincident with small luminance contrasts, as Heywood and colleagues (1994: 252) report,[19]

> [MS was] strikingly unimpaired at detecting a single coloured square concealed in a grey checkerboard when the colour is maximally saturated and the range of luminance contrasts is small. In short, saturated chromatic and achromatic boundaries are conspicuously different to M.S., particularly in dynamic displays, when they are of similar luminance contrast. Increasing the luminance contrast of the achromatic boundary renders the achromatic and chromatic boundaries perceptually similar.

Images A and B of Figure 8 are examples where MS failed and succeeded, respectively, to detect the 'different' patch in a 7x5 checkerboard display. In both cases, the edges between the target square and the surrounding squares have coincident chromatic and luminance contrasts. When the range of luminance values is within 19.4-38.48cd/m^2, as in image A, MS can detect the coloured square. Under similar conditions, MS also 'readily detected a red square concealed among greens of random, but similar luminances,' (1994: 249). When the luminance range increases to 0.01-122.2 cd/m^2, MS can no longer detect the target. Remarkably, however, MS was able to locate a coloured form, such as a cross or large square, embedded in an array of small squares varying widely in saturation and/or luminance. In image E, for example, MS 'swiftly and deftly located the concealed cross,' and in F he 'flawlessly performed a pattern discrimination between two squares and a cross,' (1994: 251).

---

[19] Thanks to Kathleen Akins for alerting me to this point.
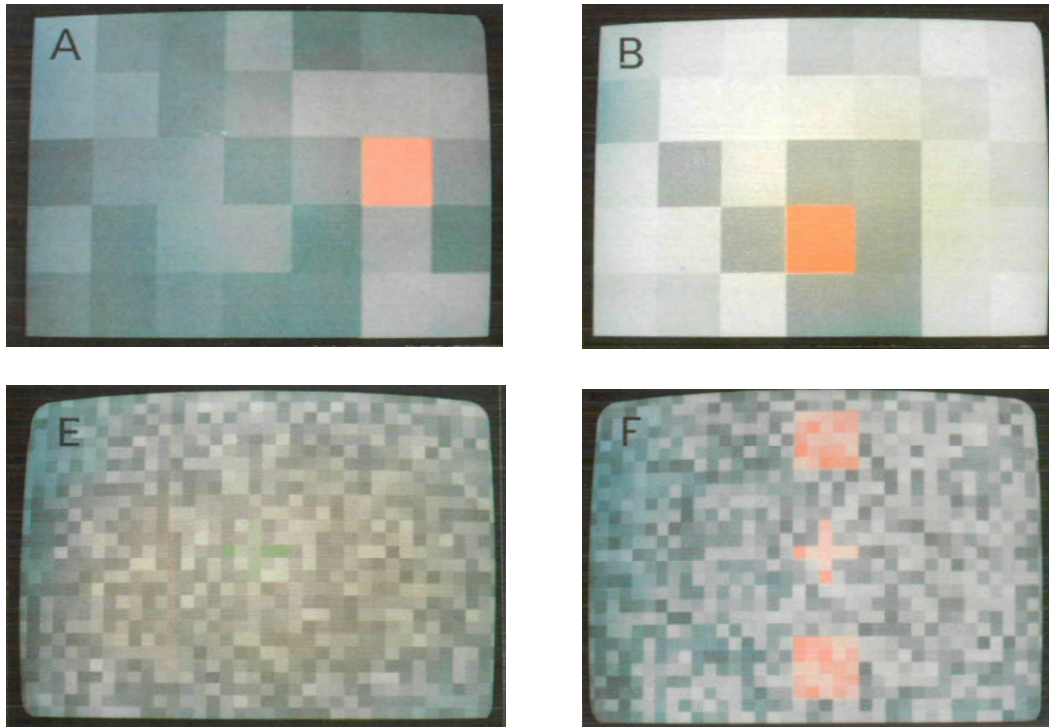
Figure 8 (reproduced from Heywood et al. 1994)

This last finding is particularly striking. When the target and distractors are of the same shape and size, as in A and B, MS can dissociate coincident chromatic and luminance contrasts, so long as the total range of luminance values is fairly small. When the target and distractors differ in shape and size, as in E and F, however, MS is able to dissociate coincident chromatic and luminance contrasts, even when the range of luminance values is quite large. It is unclear what to make of this. One hypothesis is that the 'decision' to preserve a chromatic contrast signal is not a spatially local matter. Detecting the cross in E, for example, might involve comparing the chromatic and luminance contrasts across a medium-range spatial 'window,' to determine whether there is any pattern in the chromatic contrasts that suggests the presence of a form. As a rough heuristic, if no such larger pattern is discerned, then chromatic and luminance contrast can be processed together. If a larger pattern is discerned, then it behoves form integration to keep the chromatic contrast signal separate. Preserving a chromatic contrast, then, might be a context- or task-dependent matter. In Akins' (1996: 353-4) evocative phrase, perhaps the form integration system places

'narcissistic' demands for chromatic information, as required to segment the scene. When such information is needed, it is bound with local edge representations, to facilitate the 'linking' of these segments into coherent object boundaries.[20] When it is not, perhaps the chromatic edge model does not apply.[21]

Cerebral achromatopsia forms an important part of my argument for the chromatic edge model. As such, I shall now consider three objections, and offer replies. The first objection is that MS's task performance might reflect a type of blindsight for monadic colour, rather than chromatic edge representation.[22] After all, his behaviour fits the classic blindsight pattern of preserved visual discrimination, despite severely impaired visual experience. However, blindsight patients deny having seen the stimuli, and must be encouraged to make guesses. MS needs no such encouragement. As Heywood and colleagues (1998: 413) report, 'in no case, is a patient required to "guess" the identity of an invisible figure concealed in the colour plates.' Rather, 'residual processing, of whatever origin, resulted in a conscious perceptual change, notwithstanding the absence of colour qualia,' (1998: 415). Similarly, when presented with image E in Figure 8, Heywood and colleagues (1994: 251) observed that MS 'swiftly and deftly located the concealed cross... Moreover, he did not merely detect a target; when asked he described its shape both visually and by drawing a cross with his finger.'

The second objection is that MS might have residual visual representations of monadic colours, but of a degraded or unusual sort, not describable using ordinary colour

---

[20] C.f. Kingdom et al. (1992) on chromatic 'contour linking processes.'

[21] This suggests a *pluralism* about our edge representation capacities. In Davies (2016), I appeal to context- and task-dependence in our colour constancy capacities to argue for a similarly pluralistic view of constancy.

[22] 'Blindsight' here means *type 1* blindsight, where patients supposedly have no awareness of stimuli presented in the impaired visual field, yet can discriminate some of their properties. *Type 2* blindsighters have some residual awareness of stimuli in their 'blind' field, which facilitates discrimination. This residual awareness might be severely degraded or unusual. An alternative interpretation, then, is that MS has type 2 blindsight, perhaps involving awareness of *partial* or *alien* colours. I consider this objection below.

concepts. For example, MS might represent 'partial colours,' qualities defined by at most two

of the standard three dimensions of hue, saturation, and lightness: for example, hue with

neither saturation nor lightness, or saturation and lightness without hue.[23] Alternatively, MS

might represent 'alien colours,' properties located in quality spaces quite different to our own

hue, saturation, and lightness space.[24]

My first response is that neither suggestion fits with patients' first-person

descriptions. Complete cerebral achromatopsics routinely claim that the world appears 'drab

and grey' (Cowey & Heywood, 1995: 90) or 'drained of colour' (Heywood et al., 1987: 22).

It is unclear why they would say this, if they perceive partial or alien colours. If patients

perceived partial colours, we might predict that they would describe the world as appearing

'coloured, but washed-out,' rather than 'drab and grey.' Similarly, if patients perceived alien

colours, why would they say the world looks 'drained of colour,' as opposed to 'unusually

coloured?' Both proposals seem unmotivated.

Second, there is no behavioural evidence for partial or alien colour representation. Of

course, absence of evidence is not evidence of absence. Nonetheless, if MS had partial colour

representations involving hue, for example, we would predict a capacity to identify at least

coarse-grain or determinable hue categories, such as red or green. MS has no such ability.

Interestingly, Victor and colleagues (1989) did find evidence of partial colour representation

in a patient with partial cerebral achromatopsia – so-called 'dyschromatopsia'. Their patient

identified the red bars of a grating as 'red,' and a green patch as 'tan or light green.' To date,

MS has produced no such reports. Similarly, if MS had alien colour representations, then we

would predict *some* discernible order in his groupings of colour stimuli. MS's score on the

Farnsworth-Munsell 100 hue test, however, is no better than achievable by responding

---

[23] Brown (2014). Thanks to a referee for raising this possibility.
[24] MacPherson (2015: 120).

randomly (Heywood et al. 1994: 246-7). As MacPherson (2015: 120) notes, if MS represents alien colours, it is also unclear why we can thwart his ability to complete the odd one out task by surrounding the discs with black annuli. Moreover, the alien colour hypothesis fails to explain why MS is able to detect the green and red crosses in images E and F of Figure 8, but not a single green or red square embedded within similar arrays.

My third response presents a fall-back position. Suppose I am wrong, and that MS does perceptually represent partial or alien colours. It would not immediately follow that these representations provide a superior explanation of all of MS's colour-related discriminations. Suppose MS represents partial colours with saturation and lightness dimensions, but no hue dimension. This might explain some of his discriminations, but not all, given that MS can discriminate equiluminant figures that differ in hue but not saturation. Similarly, suppose MS represents alien colours, wherein human-red stimuli and human-green stimuli do not map to opponent alien colours. Again, this might explain some of MS's discriminations, but not all, given that MS makes various discriminations indicative of red-green opponent processing. Depending on the details, then, partial or alien colour hypotheses may still leave room for explanations via the chromatic edge model.

The third and final objection is that, even if the chromatic edge explanation of MS's abilities is correct, I haven't ruled out the possibility that MS computes chromatic contrasts from very low-level monadic colour representations.[25] To paint a picture, suppose the early colour vision system is pointillist, attributing monadic colour properties to all locations in the scene. It uses these point-colours to compute chromatic contrasts, which get into edge representations in the way I describe. Later on, however, the system loses the point-colours,

---

[25] Thanks to a referee for raising this objection.

leaving only the chromatic edge representations. This is consistent with the causal MDT, if not the metaphysical MDT.

In reply, it seems extremely unlikely that early colour vision is pointillist in this sense. Vision scientists almost universally agree that low levels of the visual system code stimulus contrasts, rather than absolute values. As Hardin (1988: 27) notes, 'the common informational currency of the chromatic visual system consists entirely of output differences which are a function of… cone-excitation ratios.' As we shall see in the next sub-Section, the neural substrates for chromatic edge representation might be as early as V1 or V2. At least on a simple hierarchical, feedforward, view of colour processing, there is no reason to think that these regions code for monadic colour properties. Things are more complicated on a recurrent processing view, and I do not have space to digress here. I argue in Section 4, however, that fans of predictive coding or Bayesian architectures should still reject the MDT in favour of a *no priority* view, rather than the *relations* or *contrast first* view suggested by the simple feedforward architecture.

3.3. Neuroscience of border ownership

My view is that we perceptually represent edges holistically with geometric, topological, and chromatic features. In support of the holistic representation of geometric and chromatic features, the vast majority of cells in V1 and V2 have been found to be selective, in varying degrees, for both colour and orientation information (Friedman and colleagues, 2003: 602). Gheorghiu and Kingdom (2007: 1946) describe colour and orientation as 'conjointly represented' at this stage of visual processing. Shapley and Hawken (2011: 701) infer that 'colour and form… are inextricably linked as properties of objects in visual perception and in the visual cortex.' Gegenfurtner and Kiper (2003: 192) take such evidence to 'contradict the

idea that colour information is treated solely by a dedicated, specialized population of neurons in V1 and V2.'

My main concern is with the holistic representation of topological and chromatic features. Figure-ground phenomena suggest that we represent edges as 'belonging' to figural regions rather than grounds.[26] Zhou and colleagues (2000: 6595) present evidence from macaques, indicating that 'this perceptual tendency to assign borders to objects is reflected in the neural activity at early cortical levels.' They found evidence of cells jointly 'coding border ownership and polarity of edge contrast,' where 'polarity' means the spatial orientation of the contrast, such as whether a contrast from red to green occurs from left to right, rather than right to left, relative to the subject's position (2000: 6597). These cells pervade V1, V2, and V4. For example, around 44% of the edge-selective cells in V2 were of this type. Figure 9 shows the activity of one such cell. The peak firing rate for this cell involved stimuli with a horizontal edge between dark red and grey, with the grey 'above' the dark red, where the edge is the top edge of a dark red figure, as in C. An edge with the exact same spatial contrast polarity, but which is the bottom edge of a grey figure as in D, does not produce much response at all.[27]

---

[26] See Davies (m.s.) for further discussion.
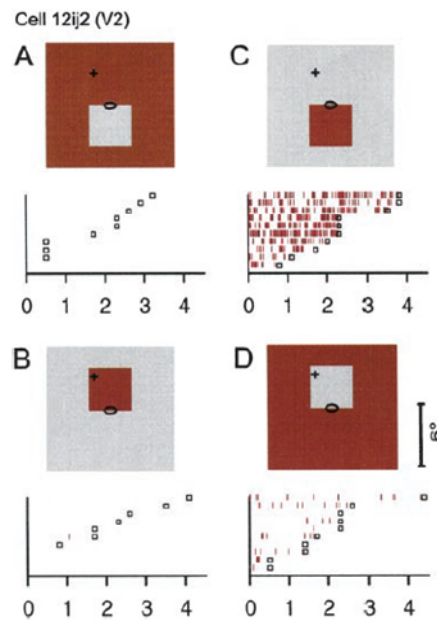[27] Friedman et al. (2003: 602) report similar results.

Figure 9 (Reprinted with permission from Zhou et al. 2000)

Clearly, we should be cautious drawing conclusions from such evidence for theories of perceptual representation. I will provide more detailed evaluation in future work. The evidence is nonetheless extremely suggestive. It seems that macaque visual cortex encodes edge information in a holistic way, with many cells responding only in the presence of stimuli that combine specific geometric (i.e. orientation), topological (i.e. border ownership), and chromatic (i.e. directed contrast) features. These cells' activity correlates closely with perception, in that border ownership signals almost always correspond with the macaque's (presumed) awareness of simple square figures (2000: 6607). The leap from here to human vision is fraught. With due caution, however, I note that the differentiation of chromatic contrast types among edge-detection cells in macaques coheres with the human behavioural evidence for contrast-dependent aspects of edge perception discussed in Sections 3.1 and 3.2. Although inconclusive, the data therefore provide confirmatory evidence for the chromatic edge model.

3.4. Function and Binding

My final argument concerns functional constraints on content attribution. I consider both input- and output-side constraints, focusing more on the latter. Regarding the input-side, the evidence from Liebmann and cerebral achromatopsia suggests that variations in the chromatic contrasts of stimuli produce variations in certain aspects of edge perception, and that such variations can be perceptually discriminated. We have also seen that certain edge-selective cells in macaque cortex respond selectively to directed chromatic contrasts. It seems unambiguous, then, that chromatic contrast is among the stimulus properties causally implicated in the production of visual edge representations. I shall not further rehash material in arguing this point.

Output-side issues prove more contentious. Although it is hard to find clear textual evidence, my impression is that many vision scientists assume that this input-side role exhausts the colour-related function of edge representation. That is, they assume that chromatic contrast information merely *feeds into* edge detection, helping the system figure out where the edges are. Clark (2000: §6.4) suggests a similar view, whereby sensitivity to chromatic contrast is useful for 'camouflage breaking,' detecting surface boundaries for the purposes of object differentiation and scene segmentation.[28] If the goal were merely to detect the *presence* of an edge, however, then presumably there would be no need for the system to use precious bandwidth retaining chromatic information to represent the *type* of edge that has been detected.[29]

Against this view, I argue that output-side functions indicate that edge representations sometimes have chromatic contrast content.[30] Regarding 'long-arm' functional role, I argued above that MS's edge perception guides his behaviour on implicit colour tasks. For example,

---

[28] C.f. Akins (2014: 200-1).

[29] Mullen et al., (2000: §4.1) consider related arguments for a single 'contour extraction mechanism' sensitive to chromatic information only at its input.

[30] 'Sometimes,' because depending on the context/task, it might not be necessary or beneficial for the visual system to represent such contrasts.

it seems that differences in MS's perception of the edges in Figure 6b enable him to select the odd one out among the three green stimuli. While MS does not describe the stimuli or their edges in chromatic terms, his behaviour is clearly sensitive to differences in chromatic contrast. This provides good reason to attribute chromatic contrast content to his visual edge representations, for these representations seemingly drive the colour-sensitive behaviour.

Regarding 'short-arm' functional role, my argument centres on the putative consumption of edge representations by processes involved in generating monadic colour representations. This is speculative, as we do not fully understand these interactions. Nonetheless, as Kentridge and colleagues note (2004b: 162), on many candidate theories, 'edge processing… [is] a starting point for surface colour perception'. Suppose that's right. And assume, as argued above, that edge processing is sensitive to chromatic contrast at its input. Now, the input-side-only view is that the edge representation process *discards* all this chromatic contrast information at its output. As Gegenfurtner (2003: 568-9) observes, however, this 'implies that the visual system would compute a colour signal for detecting edges, only to ignore it later on when determining the colour of the edge.' In terms of ideal systems design, it would make sense to *retain* this information, because chromatic contrast provides excellent information about the colours of the surfaces bounded by the edge. For example, an edge contrast of 17 units in the red direction tells us that the figure lies somewhere between 17 and the upper bound on the red dimension. Another edge contrast elsewhere of -10 units in the red direction tells us that, in fact, the figure must lie between 17 and 10-units-short of the upper bound on the red dimension.[31] If surface colour processes

---

[31] See the discussion of Morrison (forthcoming) on 'region narrowing' in the next Section. One complication is that, as noted above, the colour space for chromatic edge content is likely to differ from that for monadic colour content. Any constraints posed by the former on the latter therefore would involve translation between spaces. I will follow up on this in future work. Thanks to Nick Shea for discussion.

consume edge representations, the ideal would be to 'tag' these representations with

chromatic contrast features.

Friedman and colleagues (2003: 610) present a related argument,

[T]he factorial representation of colour and border information… appears as a

purposeful design that ensures that colour information is carried along with form

information… If colour and form information were represented in separate maps at

this level, special mechanisms would be needed to "bind" these attributes together for

a given object. However, such maps apparently do not exist in the monkey visual

cortex.

The key point is that tagging edge representations with chromatic features short-circuits the

binding problem. It is notoriously difficult to discern how, or where, the visual system binds

putatively separate representations of colour and form together in a unified representation of

a coloured surface. Friedman and colleagues suggest that that is because these representations

are not, in fact, separate: border signals carry chromatic information. This supports the

chromatic edge model, though note that my view is that edge representations have chromatic

contrast content, rather than colour content per se, as Friedman and colleagues seem to

suggest. Nevertheless, if surface colour is computed from these contrast representations, then

we still obviate the need for binding, because the contrasts are already 'bound' in a holistic

representation of the edges of the surface.

4. Alternatives to the MDT

Forthcoming in *Philosophy and Phenomenological Research*. Please cite published version.

14/01/2020

The chromatic edge model shows how a visual perceptual system might represent colour relations without representing any colours. This undermines the orthodox MDT, on which monadic colour representation determines relational colour representation. This Section surveys alternative views. Given space limitations, I remain neutral on which should be preferred.

The *relations first* view holds that visual perceptual representation of colour relations determines the visual perceptual representation of monadic colours. This inverts the standard order of explanation. An analogous view of visual spatial perception would be that representing objects as standing in certain spatial relations determines our representation of these objects as having particular spatial locations. Applied to temporal perception, representing temporal succession would be explanatorily prior to representing events as occurring at particular times. As before, we could take this as a metaphysical thesis, on which the relational grounds the monadic; or as a causal thesis, on which hypothetical mechanisms compute monadic colours from collections of relational colour representations.

The causal relations first view has not been defended to date. I plan to discuss it in future work. Morrison (2015) defends the metaphysical relations first view. In an interesting development, Morrison (forthcoming) suggests that Evie's total lifetime experience of the colour relations borne by the strawberry to other objects grounds her current visual perception of its colour. This monadic colour representation may be more or less determinate, depending on how far her prior relational perception narrows down the region of colour space into which the strawberry falls. By analogy, given two points a and b in the unit interval [0, 1], if a is 0.2 to the left of b, this determines that a is in the region [0, .8], b is in [.2, 1]. If we add the information that a third object c is 0.3 to the right of b, this narrows things down further, such that a must be in the region [0, .5], b in [.2, .7], and c in [.5, 1].

An important difference from my view is that Morrison does not endorse – or entertain – the possibility of relational colour representation without monadic colour representation. Morrison's view is that any relational colour representation, however minimal, suffices for some type of monadic colour representation. In early infancy, when Evie represents a colour relation for the first time, she will at least represent the relata as *coloured*, or perhaps *highly-determinably-red* or *highly-determinably-green*. One possible motivation for this view is that it avoids having to explain how monadic colour representation emerges from the relational at a certain level of complexity. By analogy, panpsychists avoid the thorny issue of explaining how minds emerge at some level of neural complexity, by holding that any neural state suffices for some type of mentality.[32] Another possible motivation is that the intuitive pictorial, pointillist, view of perception makes it hard to conceive what it might be like to represent colour relations without representing *any* colours.

The chromatic edge model undercuts the second motivation. It aids the positive imagination of cases in which visual perception represents colour relations between objects, despite failing to represent their colours. Regarding the first motivation, I see no problem in principle with the view that monadic colour representation emerges only given a certain complexity in relational representation. Similarly, the majority view is that some collections of neural states do not suffice for any mentality, perhaps because they are too basic, or not organised in the right way. With concessions on these points, a defender of the metaphysical relations first view could endorse my account.

A second, *no priority* view holds that there are no relations of causal or metaphysical determination between visual perceptual representation of monadic colours and colour relations. This broad view is consistent with various more specific metaphysical and causal

---

[32] Nagel (1979).

theses. On the metaphysical side, at one extreme is the view that the grounds for monadic and relational colour representation are completely distinct and non-overlapping. At the other extreme is the view that these types have precisely the same grounds. On the causal side, at one extreme is the view there is no causal interaction whatsoever between the mechanisms of monadic and relational colour representation. A moderate view is that these types emerge from complex, looping, interactions between multiple mechanisms, operating in parallel rather than in sequence. At the other extreme is the view that these types result from a single computational mechanism. This latter view might appeal to fans of hierarchical Bayesian architectures.[33] Perhaps an internally generated representation of monadic colour (the prior) is 'tested' against the sensory input, consisting of relational colour representations. A mismatch produces a prediction error signal, and the monadic colour representation is 'updated.'

Papineau (2015) endorses a no priority view, although it is unclear exactly where it falls in the above taxonomy. Papineau's view is that monadic and relational colour representations result from two distinct, doubly-dissociable, mechanisms. He postulates a 'difference detector mechanism' that issues in relational colour percepts, which represent directions and magnitudes of colour contrast. This difference detector can function independently from a 'colour classifier' mechanism, which issues categorical monadic colour percepts. As such, Papineau – unlike Morrison, but like me – allows the possibility of primitive relational colour representation. Also like me, Papineau draws on data concerning cerebral achromatopsia to support this view. The argument of Section 3.2 addresses various complications not considered by Papineau. A further substantive difference is that Papineau does not characterise 'difference detection' in terms of edge representation.[34] His view (2015: 274, italics added) is that 'detection of colour differences is a *gestalt phenomenon*: we can often consciously see straight

---

[33] Thanks to a referee for prompting me to consider this issue.
[34] Papineau (2015: 283) mentions the role of colour vision in form detection, but does not develop his account in this direction.

off that two adjacent surfaces are different in colour without first consciously having two different responses to each surface.' The notion of 'Gestalt colour differences' is not explained, and would benefit from further analysis. It might be consistent with the chromatic edge view; it might not.

5. Conclusion: Colour Relations in Form

The empiricists' analogy between visual perception and painting makes the MDT seem inevitable. One constructive project of this paper has been to argue that an iconic, quasi-pictorial, view of perceptual content does not in fact necessitate the MDT. Whereas standard views build up the representation from pixels or surface regions with monadic colour values, the chromatic edge model introduces contrast content at the ground floor, in the edge primitives: *colour relations in form*. In the limit case, this allows for the representation of chromatic contrast between two regions, in the absence of any representations of their monadic colours. The second constructive project has been to argue that this model might be true of creatures like us. Each of the empirical arguments in Section 3 warrants a paper in its own right. This paper consolidates these issues to present the strongest case for the package view. As for the preferred alternative to the MDT, my views will have to await another occasion.[35]

References

Akins, K. (1996). Of Sensory Systems and the "Aboutness" of Mental States. *Journal of Philosophy*, **93(7)**, 337-72.

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

Brown, D. (2014). Colour layering and colour constancy. *Philosophers' Imprint*, 14(15), 1-31.

Burge, T. (2010). *Origins of Objectivity*. New York: Oxford University Press.

———— (2014). Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints. *Philosophy and Phenomenological Research*, **88(2)**, 485-501.

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

Casati, R. (2003). Representational Advantages. *Proceedings of the Aristotelian Society*, 103(3), 281–298.

Casati, R. & Varzi, A. (1999). *Parts and Places: The Structures of Spatial Representation*. MIT Press.

Chirimuuta, M. & Kingdom, F. (2015). The Uses of Colour Vision: Ornamental, Practical, and Theoretical. *Minds and Machines*, 25, 213-29.

Clark, A. (2000). *A Theory of Sentience*. OUP.

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

Cowey, A., & Heywood, C. A. (1995). There's more to colour than meets the eye. *Behavioural Brian Research*, **71**, 89-100.

Davies, W. (2014). The inscrutability of colour similarity. *Philosophical Studies*, 171(2), 289-311.

———— (2016). Color constancy, illumination, and matching. *Philosophy of Science*, 83(4), 540-62.

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

———— (m.s.). Visual border disputes: Figuring out figure-ground.

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

Gegenfurtner, K. R. (2003). Cortical mechanisms of colour vision. *Nature Reviews Neuroscience*, 4, 563-72.

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

Gheorghiu, E., & Kingdom, F. A. A. (2007). Chromatic tuning of contour-shape mechanisms revealed through the shape-frequency and shape-amplitude after-effects. *Vision Research*, **47**, 1935-49.

Hardin, C. L. (1988). *Colour for Philosophers: Unweaving the Rainbow*. Hackett Publishing Co.

Heywood, C. A., Cowey, A., & Newcombe, F. (1991). Chromatic Discrimination in a Cortically Colour Blind Observer. *European Journal of Neuroscience*, **3**, 802-12.

——— (1994). On the role of parvocellular (P) and magnocellular (M) pathways in cerebral achromatopsia. *Brain*, **117**, 245-54.

Heywood, C. A., Kentridge, R. W., & Cowey, A. (1998). Form and motion from colour in cerebral achromatopsia. *Experimental Brain Research*, **123**, 145-53.

Heywood, C. A., Wilson, B., & Cowey, A. (1987). A case study of cortical colour "blindness" with relatively intact achromatic discrimination. *Journal of Neurology, Neurosurgery, and Psychiatry*, **50**, 22-9.

Hume, D. (2000). *A Treatise of Human Nature*, D. F. Norton & M. J. Norton (eds.), OUP.

Kentridge, R. W., Heywood, C., & Cowey, A. (2004a). Chromatic edges, surfaces and constancies in cerebral achromatopsia. *Neuropsychologia*, **42**, 821-830.

Kentridge, R. W., Cole, G. G., & Heywood, C., (2004b). The primacy of chromatic edge processing in normal and cerebrally achromatopsic patients. *Progress in Brain Research*, **144**, 161-9.

Kingdom, F., Moulden, B., & Collyer, S. (1992). A Comparison Between Colour and Luminance Contrast in a Spatial Linking Task. *Vision Research*, 32(4), 709-17.

Koffka, K. (1936). *Principles of Gestalt Psychology*. Kegan Paul, Trench, Trubner & Co. Ltd.

Kosslyn, S., Thompson, W., & Ganis, G. (2006). *The Case for Mental Imagery.* OUP.

Livingstone, M. & Hubel, D. (1987). Psychophysical Evidence for Separate Channels for the Perception of Form, Color, Movement, and Depth. *The Journal of Neuroscience*, 7(11), 3416-68.

Locke, J. (1975). *An Essay Concerning Human Understanding*, Peter H. Nidditch (ed.), OUP.

MacPherson, F. (2015). The structure of experience, the nature of the visual, and type 2 blindsight. *Consciousness and Cognition*, **32**, 104-28.

Mandik, P. (2014). 'What Is Visual and Phenomenal but Concerns Neither Hue Nor Shade?' In Brown (ed.), *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience*, (pp. 225-33). Springer Press.

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

Morrison, J. (2015). Anti-Atomism about Colour Representation. *Nous*, **49(1)**, 94-122.

——— (Forthcoming): 'Perceptual Variation and Structuralism,' *Nous*, early view: **https://doi.org/10.1111/nous.12245**.

Mullen, K. et al. (2000). Contour integration in color vision: a common process for the blue–yellow, red–green and luminance mechanisms? *Vision Research*, 40, 639-55.

Nagel, T. (1979). 'Panpsychism', in *Mortal Questions*, Cambridge: Cambridge University Press.

Papineau, D. (2015). 'Can we really see a million colours?', in Coates and Coleman (eds.), *Phenomenal Qualities: Sense, Perception, and Consciousness*, (pp. 274-). OUP.

Pautz, A. (2006). Can the Physicalist explain colour structure in terms of colour experience? *Australasian Journal of Philosophy*, **84(4)**, 535-64.

Quilty-Dunn, J. (2017). *Syntax and Semantics of Perceptual Representation*, CUNY Academic Works. http://academicworks.cuny.edu/gc_etds/2261

Shapley, R., & Hawken, M. (2011). Color in the Cortex: Single- and Double-Opponent Cells. *Vision Research*, 51, 701–17.

Shevell, S. & Kingdom, F. (2008). Color in Complex Scenes. *Annual Review of Psychology*, 59, 143-66.

Smith, B. (1998/99). Boundaries: A Brentanian Theory. *Brentano Studien*, 8, 107-14.

Victor, J. et al. (1992). Acquired central dyschromatopsia: Analysis of a case with preservation of colour discrimination. *Clinical Vision Science*, 4(3), 183-96.

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

Zhou, H., Friedman, H., & von der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. *Journal of Neuroscience*, **20**, 6594–6611.