King's Research Portal

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Papineau, D. (Accepted/In press). *Can We Really See A Million Colours?* (pp. N/A). Kings College London, University of London. <http://philpapers.org/rec/PAPCWR-2>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Can We Really See a Million Colours?¹

David Papineau

Abstract

Nearly all philosophers assume that human beings are capable of well over a million different conscious visual responses to coloured surfaces (and they then debate whether this shows that some mental representation is non-conceptual). I deny the premise. In this paper I argue that humans are capable of only a limited range of colour responses to surfaces presented singly (along with a gestalt ability to register that adjacently presented surfaces are different-in-colour). This may be counterintuitive, but it accommodates the empirical data better than the standard view, and also resolves various philosophical puzzles.

Contents

1 Introduction

2 Two Kinds of State

3 The Debate about Non-Conceptual Content

4 Picture and Features

5 Functions of Colour Perception

6 Mechanisms and Dissociations

7 The Intransitivity of Indiscriminability

8 Intransitivity Resolved

9 Conclusion

1. Introduction. It is widely assumed that human beings are capable of well over a million different conscious visual responses to coloured surfaces.

The evidence for this is the psychometric data on ‘just noticeable colour differences’. When two coloured surfaces are presented simultaneously and adjacently, normal human beings can detect very fine differences. They are capable of consciously registering colour differences between well over a million differently coloured sample surfaces, varying in hue, brightness and saturation.²

It might seem natural to infer from this that we humans must have a repertoire of more than a million conscious responses to coloured surfaces. After all, when we are presented with a pair of adjacent surfaces, surely our consciously registering a difference in colour must derive from our first having one colour response to the left hand side surface, and another colour response to the right hand surface, and thence registering that there is a difference.

However, this assumption, which I shall call the ‘orthodox view’, is not mandatory. In this paper I shall offer an alternative account of the discrimination data. On this alternative view, detection of colour differences is a gestalt phenomenon: we can often consciously see straight off that two adjacent surfaces are different in colour without first consciously having two different responses to each surface.

This alternative view is consistent with our being capable of far fewer than a million conscious colour responses when we are presented with colour samples on their own. Perhaps we have a limited number of such non-comparative conscious colour responses. If the detection of colour differences

¹ Forthcoming in Coates, P. and Coleman, S. (eds) The Nature of Phenomenal Qualities Oxford: Oxford University Press 2014. Pre-print version archived with permission.

² Pointer and Attridge 1998, McCamy 1998. It is relevant to what follows that the ability to discern simultaneous colour differences between non-juxtaposed surfaces falls off rapidly with increasing spatial separation (Sivic 1997 163).

between adjacent surfaces does not derive from prior non-comparative responses to each surface, there is no need to posit a million such responses to account for the discrimination data.

The alternative view fits well with the fact that humans perform very poorly on same/different tasks when colour samples are presented successively rather than simultaneously. When two coloured surfaces are presented at separate times, the number of different surfaces that can reliably be judged to be the same or different by most observers is of the order of a few dozen, even under ideal conditions (Hamwi and Landis 1955).

Orthodoxy attributes this poor performance in the diachronic successive matching task to some kind of memory limitation. On the orthodox view, subjects are capable of distinct conscious responses to a great many more than a few dozen different coloured surfaces; so their inability to make reliable same/different responses to more than a few dozen successively presented surfaces must be due to some sort of failure to retain knowledge of their earlier response until the later presentation.

But, on my alternative view, we don't necessarily have to attribute the poor performance in the successive matching task to memory deficiency. Perhaps most humans are capable of no more than a few dozen conscious responses to coloured surfaces presented singly in the first place, and the reason that they aren't reliably attuned to finer distinctions in the successive same/different task isn't that they can hold the relevant fine-grained information in memory, but simply that they didn't have it in the first place.

However, this is more than I need, or indeed want, to say. The main burden of my alternative view is that we have far fewer conscious colour responses to surfaces presented singly than the million plus assumed by orthodoxy. It would be an extra claim to equate this number with the few dozen categories over which subjects can perform reliably in the successive same/different tasks. Maybe subjects do have significantly more than a few dozen available conscious responses to surfaces presented singly, but fail to manifest these in the successive matching tasks because of difficulties in retaining the relevant information in memory.³

As will become clear in what follows, my alternative view makes it a matter of internal cognitive architecture how many different conscious colour responses humans can make to surfaces presented singly. The number of such responses depends on the organization of the brain, and is not something that can easily be ascertained from psychometric or even introspective data.

We should also bear in mind that the range of such responses is very likely to vary from individual to individual. It will be entirely consistent with my alternative view that variations in culture, training and natural endowment should make a significant difference to the repertoire of non-comparative colour responses available to different individuals. Maybe some individuals are indeed only capable of a few dozen such responses, while others—painters or interior decorators, say—are capable of many hundreds.

For both these reasons, it would make no sense for me to offer any definite number for the number of colour responses humans in general are capable of. Still, having made these points, it will sometimes be useful to have a number to play with, and so I shall on occasion assume that there are a hundred single-colour responses available to humans. But this will be for illustrative purposes only, and

³ This diagnosis receives *prima facie* support from evidence that performance in colour memory tasks is enhanced if the subjects have names for the colours involved (Roberson et al 2005). Of course this evidence does not prove that the limited performance is due to memory constraints—the names may matter, not because they help subjects to remember colours, but because they shape the single-colour responses which subjects are capable of in the first place (cf Roberson et al 2004, Winawer et al 2007). It is consistent with my alternative view that subjects' single-colour responses cannot outstrip the colour categorizations made by their culture and reflected in their colour vocabulary. But, again, this is more than I need or want to argue here. For a survey of relevant recent data in this hotly debated area see Regier and Kay 2009.

should not be taken to deny that the real number will depend on details of brain structure, and moreover will vary across individuals.

More generally, in what follows I shall be offering an alternative to the orthodox account of the structure of colour perception. However, this alternative will be something of a speculative simplification. No doubt a full understanding of colour perception will show it to be more complicated than I portray it here. The physiology of colour vision is by no means fully understood, and future findings are sure to add to our understanding of the structure of conscious colour vision.

However, the primary purpose in this paper is not to establish a precise positive account of conscious colour vision, but to cast doubt on orthodoxy. For this purpose, it will be convenient to formulate a simple alternative, and to show that the existing evidence favours it over orthodoxy. But it should be borne in mind that variations on this alternative would be equally well-supported by that existing evidence, and may no doubt be favoured by future evidence.

One final preliminary. My focus in all that follows will be on phenomenal properties. That is, I shall be concerned with the range of ways it can consciously be for an individual in colour vision. In this respect, I shall have a narrower focus than much other philosophical work on colour vision, in that I shall aim to bracket off questions about the ontology of colour. I shall be exclusively concerned with the range of conscious mental states involved in colour vision, and not with the range of properties of surfaces that such states might be tracking.⁴

Moreover, I shall aim to remain as neutral as possible about the nature of conscious phenomenal properties. As it happens, I think that conscious vision is designed to represent features of the external environment, and so will on occasion talk in about conscious colour perception in representational terms; I also think that conscious phenomenal colour properties are to be identified with ‘syntactic’ features of the vehicles of such representation, rather than with the abstract representational contents of those vehicles. However, I do not think that the main points in what follows depend on these commitments. Those who have other views about the nature of the phenomenal properties should be able to transpose my arguments into their own theoretical frameworks.

2. Two Kinds of State. On the view I propose, our conscious colour experience is the joint product of two different kinds of perceptual state.

The first kind of state arises in response to single surfaces, and is designed to categorize the surface as of a certain colour. Here I hypothesize that the visual system is capable of something like one hundred different colour responses. I shall refer to perceptual states of this kind as ‘categorical’. We can think of them of corresponding to ordinary colour terms, such as pink, orange, purple, navy blue, olive green, and so on.⁵

⁴ Pete Mandik ‘Color-consciousness conceptualism’ (2012) appeals to a distinction between non-comparative and comparative colour perceptions that is close to that defended in this paper. But his primary focus is on the representational contents of the perceptions rather than their qualitative nature, and it is not clear to me whether he ends up endorsing my thesis that humans are capable of far fewer than a million colour responses. Mandik acknowledges David Rosenthal’s important earlier discussion of non-comparative and comparative colour concepts (2005 188-9). But Rosenthal too seems not to deny that humans are capable of a million-plus ‘mental colours’. The relation between Rosenthal’s and Mandik’s positions and my own is complicated by the fact that Rosenthal, and possibly Mandik too, will not count the million-plus different responses as conscious if they are not conceptualized, whereas I am not assuming any kind of higher-order approach to consciousness. Still, even if we put this difference to one side, I deny, where they seem to accept, that our sensory systems generate a million-plus colour responses, whether or not we count these responses as conscious.

⁵ A central piece of evidence for the now widely-accepted ‘opponent-process’ account of colour perception (Hurvich and Jameson 1957) is the ability of subjects reliably to estimate the proportion of red/green or blue/yellow in any spectral colour (Werner and Wooten 1979). This quantitative ability may seem in tension with my hypothesis of a limited range of categorical single-colour responses. However, this tension is only

The second kind of perceptual state arises when the visual system is presented simultaneously with two coloured surfaces side-by-side. This kind of state will convey information about the relation between the two surfaces. The simplest version of such a perceptual state would simply indicate whether the two samples are the same or different in colour. But it is consistent with my view that such states should also indicate the relative locations of the two samples in various colour dimensions: thus they might indicate which sample is more red than the other, and similarly for more yellow, more saturated, and brighter⁶. I shall refer to states of this kind as ‘relational’.

The crucial point for my purposes is that the visual system should be able to issue in a relational judgement that two adjacent samples viewed simultaneously differ even in cases where the two surfaces produce the same categorical colour response. That is, the two surfaces both elicit the same version of the first kind of categorical state—they both strike the viewer as orange, say—but nevertheless elicit a relational state that indicates a colour difference between the two samples.

If the relational state simply conveyed symmetrical same/different information, then this would imply that in such a case the conscious experience of the viewer would be essentially the same with respect to the two adjacent surfaces: the viewer would see the left hand sample as orange and as different from the right hand sample, and similarly see the right hand sample as orange and as different from the left hand sample.

But if the relational states also convey information about which sample is more red/yellow/saturated/bright, then the situation need not be so symmetrical. The viewer may not only judge relationally that two oranges are different, but also that the left hand sample is more red than the other, and similarly for the other dimensions. To this extent, the observer’s conscious experience of the two samples will be different: one will be seen as more red, say, and the other correspondingly as less red.

Still, my view does have this striking consequence: looking at orange₂₃ next to orange₂₄ can be just like looking at orange₂₇ next to orange₂₈. (I am here using orange₂₃, orange₂₄, and so on to refer to the surface properties of samples that are just noticeably different in colour when presented side-by-side.) In both cases—orange₂₃ next to orange₂₄, and orange₂₇ next to orange₂₈—the viewer would have the categorical information that the two adjacent surfaces are both orange, plus relational information to the effect that the left hand sample is more/less red/yellow/saturated/bright than the other. If our conscious awareness is the product of the two states conveying this information, as I am supposing, then our conscious awareness of the two pairs would be just the same.

This is in striking contrast to orthodoxy. On the orthodox view, orange₂₃ next to orange₂₄ must look different to orange₂₇ next to orange₂₈. For orthodoxy takes our awareness of colour differences to derive from our awareness of the individual colour of each sample, which implies that each orange_k must produce a different colour response on its own, and thus that different pairs of such surfaces will do so too.

Here is a way of conveying my view of colour perception. Imagine that a friend is looking at two side-by-side samples that are invisible to you. Your friend tells you, say, that the samples are both orange, and that the left hand sample is more red, less yellow, less saturated and more bright than the

apparent. The subjects’ estimates are by no means so precise as to entail that they are capable of more than a few dozen distinct categorical colour responses. And my view can perfectly well allow that these categorial responses depend on the activity of opponent channels in ways that can to some degree be recovered by introspective reflection.

⁶ This particular set of dimensions is suggested the opponent-process theory. But my view is not committed to this theory, and could equally well accommodate different dimensions of colour relations, and also more specific qualitative comparative judgements (for example: much redder/redder/same redness). I shall return to the issue of dimensions of colour relations towards the end of section 6.

right hand sample. I say that the information that your friend has provided exhausts the information that you would get even if you were looking at the samples yourself.

3. The Debate about Non-Conceptual Content. Many readers will have been introduced to questions about the range of colour responses available to humans in connection with the debate about ‘non-conceptual content’. While this debate is not my focus in this paper, it may be helpful to explain what my view implies about it.

The debate about non-conceptual content concerns the representational powers of perception. Can perception represent facts that cannot be conceptualized by the perceiving subject?

To keep things simple, let me assume that representational contents are ‘Russellian’, depending on the object and properties represented, and not on their modes of presentation. Now assume that some subject perceptually represents that some a is F. Must the subject possess concepts that refer to a and to F? The conceptualists say ‘yes’, the non-conceptualists ‘no’. (To bring out the possibility of non-conceptual representation, imagine that some insect represents that the leaf in front of it is edible, say, in a way that does not require that it can refer to that leaf, or to the property of edibility, in other representational contexts.)

One of the most familiar arguments for non-conceptual representation appeals to the ‘fineness of grain’ of colour perception. The idea is that the psychometrical data on just noticeable differences show that we are capable of visually representing over a million different colours. But it seems highly implausible to suppose that we have over a million colour concepts. (If we can conceptualise all those colours, how come we are so hopeless at retaining visual information about colours over time?) So conscious vision can represent properties for which we have no concepts. (Evans 1982, Peacocke 1992.)

Nearly all the literature on this argument accepts that the discrimination data do show that we can visually represent a million different colours. The debate has focused rather on whether or not we have a corresponding range of colour concepts.

Conceptualists have responded to the ‘fineness of grain’ argument by appealing to ‘demonstrative’ colour concepts. The idea is that such ‘demonstrative’ colour concepts can be formed when a specific colour sample is in view, and can be expressed by the phrase ‘that colour’. Since there is nothing to stop us forming concepts like these for all the million-plus colours we can represent, so the conceptualists argue, there is no reason to suppose that our representational powers outrun our conceptual repertoire. (McDowell 1994, Brewer 1999.)

Non-conceptualists have objected that these putative demonstrative concepts do not really qualify as concepts. A minimum requirement for the possession of concepts, they argue, is that such possession subserves the retention of information over time. But subjects are typically unable to retain information about fine-grained colours they saw a moment ago. So we should not credit them with any corresponding concepts. (Raffman 1995, Kelly 2001.)

One possibility open to conceptualists at this point is to respond that their opponents are setting the standards for concept possession too high, and that there are good reasons for crediting us with fine-grained demonstrative concepts despite their failure to allow information retention. (Brewer 2005, Chuard 2006.)

However the position I am defending in this paper allows a quite different defence of conceptualism, one that does not call for any dilution of standards for concept possession. This is because my view rejects the assumption that gets this whole debate off the ground, namely, the assumption that humans are able visually to represent a million different colours. On my view, we can only represent a hundred-odd categorial colour properties in the first place (as well as being able to represent that two surfaces seen side-by-side differ from each other in certain dimensions). So there is no need to posit a

million concepts to render all representation conceptual. A hundred or so will do (plus a few concepts to represent two adjacent surfaces being relationally different in certain dimensions).

Put it like this. My view agrees with conventional non-conceptualists that we have a highly limited number of colour concepts. But, from my perspective, this is quite enough to render all colour representation conceptual, since we don't represent any more than a highly limited number of colours to start with.⁷

4. Pictures and Features. Some readers may be feeling that this would be a fine way to defend conceptualism, if my view were true—but that as yet I have offered no earthly reason to suppose it is true. Moreover, they may well feel that the orthodox theory of a million different conscious colour responses is far more in line with obvious facts than my outlandish suggestion that we are only capable of a hundred responses or so.

In this section I shall aim to cast doubt on the idea that the orthodox view is forced on us by any obvious facts. In subsequent sections I shall offer further positive considerations in favour of my alternative view.

If conscious vision worked in a pixel-like pictorial manner, by first consciously determining the local qualities of each minimal region in the visual field, and only thence inferring the properties of larger regions, then I agree that the orthodox view of colour perception would follow. However, it is highly doubtful that conscious vision does work in this pictorial manner.

Note that the issue is not whether viewing a real scene produces the same kind of conscious response as looking at a picture. This is of course trivially true, given that pictures (at least traditional representational pictures) are designed to produce the same kind of conscious visual reactions as normal scenes. Rather the question is whether our conscious awareness of an overall scene is metaphysically determined by our prior conscious awareness of each minimal region of the field, in the way that the large-scale visible properties of a picture are arguably metaphysically determined by the visible properties of each of its minimal regions.

While it might seem natural to suppose that conscious vision works in this pictorial way, this doesn't fit what the brain does. Perhaps the information registered by the retina at any time can be considered to comprise a kind of pictorial representation of what is seen (though it will be both fleeting and highly degraded outside the central foveal region). But I take it that this retinal state is certainly not conscious. And all later stages of visual processing work in a highly non-pictorial manner. Far from trying to identify local properties of minimal regions, different processes strive to discern spatiotemporally non-local features, such as orientations, shapes, motions, and colours, prior to 'binding' together features that are co-instantiated in the same objects. The processes that discern these features progressively discard much of the information that provides input for their analysis. The whole point of these processes is to sift through noisy, multifarious and often equivocal inputs and reach a verdict on, say, the orientation of some edge or the shape of some surface. They will then transmit this verdict to other parts of the brain, but not the basis on which they reached it.

Now, it is a matter of some controversy, to put it mildly, exactly which brain states correspond to the phenomenal properties that constitute conscious vision. But, whatever the precise answer to this question, there is no immediate reason to suppose that there is a range of such possible conscious states corresponding to all the million-plus surface colour properties that are just noticeably different. It is an interesting question whether the inputs to the processes giving rise to conscious colour states contain enough information to identity all these different surface properties. But, even if they do, it does not follow that the conscious outputs of these processes include distinct states for each of these million-plus surface properties.

⁷ For a more detailed defence of conceptualism along very similar lines, see Mandik 2012.

On my view, there is one cognitive process that categorizes single surfaces into one of a hundred-odd kinds, and another that discerns a limited range of relations between adjacent surfaces. While both will use fine-grained wavelength information as input, this information will have been discarded before they produce their outputs. So on my view conscious visual perception will not be capable of anything like a million potential responses to coloured surfaces.

Of course, I have not yet given any definite argument that there isn't a cognitive process capable of producing such fine-grained conscious colour outputs. Such arguments will come in a moment. This section's purpose, remember, is only to show that there is nothing in the basic workings of visual perception to force us to recognize such a mechanism.

Does not introspection show us directly that conscious vision works pictorially? When we reflect on the nature of our visual experience, can't we just see that each minimal area of the visual field plays its part in determining the character of the whole? Well, no doubt introspection shows us that each minimal area of the visual field has a range of definite properties. But I see no reason to accept that it shows us that these properties are determined in a bottom-up way, with conscious vision first fixing the local qualities of minimal regions, and then allowing this to determine the conscious properties of the whole.

We should ask ourselves—what would it seem like if vision worked in a gestalt manner, with many relational non-local features, such as shapes and motions and colour differences, entering directly into conscious perception, independently of any conscious registration of local qualities that might determine them? I say that things would then seem to us just as conscious vision actually does. There is nothing in introspection itself to show us that conscious vision works in a bottom-up pictorial way. (Indeed, if you ask me, the idea that conscious vision works pictorially derives from confused reflection on the physical nature of pictures, rather than from any deliverances of introspection.)

While we are on introspection, it is worth remarking that there are two oft-remarked features of introspectible visual experience that fit rather better with my view than orthodoxy.

The first is the subjective categoricity of colour perception. When we look at a picture of some region of colour space, we do not experience a continuous variation in the relevant dimensions, as we do when we are noting the loudness and pitch of noises, say, or the warmth and hardness of surfaces. Rather, we see a set of distinct colour regions (orange, yellow, green, turquoise, . . .), divided by fairly definite boundaries. This argues that the conscious brain is placing viewed surfaces into a limited number of discrete categories, rather than locating them at some position along a continuous spectrum.

The second relevant feature is our ability simultaneously to see adjacent regions of some variably illuminated surface as of the same categorial colour yet as differing in colour. (Cohen 2008.) From an orthodox point of view, this can seem puzzling. When we see the whole surface as having the same colour, it's not just that we judge the different regions to be different shades of some determinable colour—different shades of red, say. Rather we see the whole surface to be just one determinate colour, one specific shade of red (our colour perception shows us that we'd need only one kind of paint to produce a surface like that). But how can this be part of our visual experience, on the orthodox view, if we simultaneously see different regions of that surface to be different colours?

However, from my alternative point of view this phenomenon is just what we should expect. The different regions of the surface are being categorized as the same colour (courtesy of the process that makes such categorial judgments), while at the same time adjacent regions are being experienced as different in colour (courtesy of the process that detects relational colour differences).

5. Functions of Colour Perception. It will be helpful to think about the functions that colour perception is designed to serve. We can distinguish two basic such functions.

First, there is the classification and reidentification of objects. Is that a robin? Is this cherry ripe enough to eat? Is that John's dog? Is this my scarf?

This first function could be served perfectly well by a repertoire of about a hundred colour categories. Anything more would seem like overkill. After all, the different instances of a given natural type—such as different robins or different ripe cherries—are likely to vary in more fine-grained chromatic respects. And even a given object—such as John's dog or my scarf—will vary in colour over time. For purposes of classification and reidentification, a million different colour categories would seem quite unnecessary, and indeed liable to lead to distinctions that are positively counter-productive.

The second function is the discernment of three-dimensional physical form. One of the primary functions of human vision is to construct a representation of the medium-sized physical objects in our environment, including their sizes, shapes, orientations, relative positions, and so on. The brain uses many clues to do this, but crucial first steps are the identification of the visible edges which mark the boundaries of objects in the visual field.

The standard assumption is that the visual brain identifies these features by detecting loci which mark light intensity differences—that is, lines where the light intensity on one side is different from that on the other. But, as we shall see in the next section, it also uses loci of chromatic wavelength differences. From a functional perspective this is unsurprising, given that loci of chromatic differences are often better guides to the presence of edges than intensity differences alone. This is because they are not generated by mere shadows in the way that intensity differences are: the light from either side of the line produced by a shadow will necessarily vary in intensity, but will generally not vary in wavelength. This point is all the more significant if we consider our evolutionary history. As Israel Abramov has observed in this connection, “primates developed in an arboreal environment, which is characterized by a bewildering and random array of leaves and shadows” (1997, 110).

Edge-detection will benefit from a very fine-grained sensitivity to wavelength differences, even if categorical colour classification does not. Any consistent locus of discontinuity in the wavelength of reflected light is likely to indicate an edge, however small. The edge-detecting function will thus be served by the maximal sensitivity that the eye and brain can manage. There is no overkill in detecting the most precise colour differences, even if we have no use for fine-grained absolute colour classifications.

Functional considerations thus give us every reason to suppose that there are two distinguishable colour-related processes in the visual brain. One will be concerned to categorize single surfaces into a hundred-odd kinds, as an aid to the classification and reidentification of objects. The other will be concerned to detect colour relations between adjacent regions of the visual field, in the service of the discernment of three-dimensional visual form.

We might model the first process as a kind of neural net with a hundred-odd output nodes. When it is presented with any surface, one of these output nodes will be activated.

We can expect this ‘surface classifier’, as I shall call it henceforth, to manifest ‘colour constancy’. That is, it will be able to judge surfaces as retaining the same colour even across a wide range of changes of illumination. Such changes of illumination can radically alter the amounts of different wavelengths of light that surfaces transmit to the eye. But such changes of illumination will of course not alter the intrinsic natures or identities of the objects with those surfaces. Given its function, the surface classifier will thus need to factor out the optical differences caused by changes in illumination, and instead track surface properties that are invariant with respect to such changes.

The second kind of process might be modelled as comparing the light received from adjacent regions in the visual field in various dimensions, and subtracting the relevant quantities from each other. It will note whether the answer is positive, negative, or zero, but will throw away the absolute values of these quantities, since these absolute values are irrelevant to its function of discerning three-

dimensional form. It will register fine-grained differences between surfaces even when these differences are too small to make any difference to the verdicts of the surface classifier.

Note how the outputs of this ‘difference detector’, as I shall call it, will also be insensitive to changes of illumination. Such changes will affect the absolute values of the optical quantities the detector is comparing, but in general will leave the relative values unchanged. This insensitivity as it should be, given that changes in illumination do not normally signal changes in three-dimensional form.

In my view, our conscious colour experience is the joint product of the outputs of these two different processes. In the first place, we experience surfaces as having one of a hundred-odd colours. But alongside this we also experience adjacent regions of the visual field as differing in colour, even in cases where the two regions prompt the same response from the colour classifier.

6. Mechanisms and Dissociations. The existence of two distinguishable colour mechanisms is amply confirmed by what is known of the neurophysiology of colour vision.

It is widely accepted that chromatic wavelength information is carried to the primary visual cortex by the lateral geniculate nucleus by two kinds of cells. There are the L/M cells that are stimulated by long-wavelength L cones and inhibited by medium-wavelength M cones (or vice versa)—these are standardly called red/green opponent cells. And then there are the S/(L+M) cells that are stimulated by short-wavelength S cones and inhibited by the sum of L and M cone signals (or vice versa)—these are standardly called yellow/blue opponent cells. (Abramov 1997 101-5, Shapley and Hawken 2011 704.)

It is also widely supposed that these two types of cells bear some relation to the opponent ‘red-green’ and ‘yellow-blue channels’ posited by the ‘opponent-process theory of colour vision’. According to this theory, experienced categorial hues are determined by the levels of activation of two such opponent channels: that is, we can think of such hues as points in a two-dimensional space whose axes are the activation levels of these two channels. This theory thus offers a nice explanation of why we can have hues that strike us as a mixture of red and blue or yellow, or green and blue or yellow, but none which seem a mixture of red and green, or of blue and yellow.

How exactly the two kinds of chromatic cells in the lateral geniculate nucleus might relate to the two posited channels of the opponent process theory is not a straightforward matter. (Hardin 1993 54-8, Abramov 1997 107). However that may be, though, it is clear that the determination of experienced categorial hues is not the only upshot of the chromatic information carried to the visual cortex by the L/M and S/(L+M) cells.

Many of the cells in the primary visual cortex are ‘double-opponent’ cells. These are cells that fire when they get contrasting inputs from opponent cells which themselves respond to different but adjacent regions of the visual field. For example, we might have a cell that is excited when L/M inputs responsive to one region of the visual field are opposite in sign to L/M inputs from an adjacent region.

Cells of this kind are designed to detect spatial changes in optical properties. They fire when adjacent regions of the visual field produce differing optical responses. Many such double-opponent cells are also orientation-selective. They fire only when the locus of optical discontinuity lies in a specific direction. The function of these double-opponent cells is clearly to serve the recovery of three-dimensional form, by helping to identify whether edges are present in some region, and determining their orientation when they are. (Shapley and Hawkin 2011 704.)⁸

⁸ If these cells are specifically sensitive to spatial variation, then they will be unable to account for any ability to detect fine-grained colour differences in samples displayed in immediate temporal succession. However it seems likely the mechanisms underlying the perception of motion involve cells that are sensitive to temporal

Until relatively recently it was supposed that the detection of form depended predominantly on achromatic information about spatial discontinuities in light intensity. However, it is now clear that the double-opponent cells in V1 that subserve the recovery of three-dimensional form are generally just as responsive to chromatic information as to achromatic information (Shapley and Hawkin 2011 707). This is unsurprising, given the point made in the last section about the importance of chromatic discontinuities for identifying edges in an environment containing shadows.

It is sometimes said that double-opponent chromatic cells provide the basis for colour constancy computations (Kentridge, Heywood and Davidoff 2003, Heywood and Kentridge 2003). But this can be a bit confusing. Certainly these cells can detect whether or not transmitted wavelength alters from one spatial region to another, and to this extent can sometimes detect that a certain area is uniform with respect to wavelength properties. And this can then inform a visual verdict that the relevant area is part of a single surface with a uniform surface colour property. Moreover, none of this processing in going to be affected by variation in background illumination, since double-opponent chromatic cells, like double-opponent cells in general, are sensitive to differences in optical properties between different regions, and not to the absolute value of these properties.

But, for just this reason, double-opponent chromatic cells on their own are no good for computing which uniform colour a surface has. The information they convey can help to determine that a given area is all one colour, but it does not determine what that colour is. The information about absolute wavelengths that is needed to identify categorial colour has been discarded by the time we get to the outputs of double-opponent chromatic cells. These cells can convey that one region is ‘more red’ than an adjacent region, say, but this information is consistent with the two regions have a wide range of categorial colours. Double-opponent cells are thus designed to identify colour changes, so to speak, not the colours themselves. So they provide a basis for colour constancy only in the sense that they help discern the three-dimensional form which is an essential precondition for the attribution of uniform colours to surfaces. They do not determine those colours themselves.

I thus take the double-opponent cells in the primary visual cortex to be part of the difference-detecting mechanism postulated in the last section. They inform a process that eventually issues in conscious judgements as to whether adjacent regions differ or coincide in colour, possibly augmented by judgements as to the relative positions of the surfaces in a few colour dimensions, such as red-green, yellow-blue, or bright-dark. This process does not reach any conclusion about the categorial colours of surfaces.

What now about the other mechanism postulated in the last section, the ‘surface classifier’? Well, the location of a ‘colour centre’ in the human brain has been a topic of intense debate over the past few decades. In the 1970s Semir Zeki showed that cells in area V4 in the monkey brain responded reliably to surface colours across changes in illumination, and to this extent seemed to be operating as a surface classifier⁹. Subsequent work has sought to identify an analogue of V4 in humans, but specific suggestions have proved highly controversial. In any case, it is not clear that V4 is dedicated specifically to colour rather than form, even in monkeys¹⁰. Moreover, there is evidence that certain areas yet further forward in human brain (‘VO-1’ and ‘VO-2’) are highly activated by coloured stimuli¹¹. A recent survey article on colour vision suggests that “perhaps VO-1 is a memory bank of colours that is compared with incoming color signals from V1” as an aid to object recognition using “association of current input from the world with prior experience”¹².

variation in chromatic information (Shapley and Hawken 2011 703-4). If so, this would explain an ability to detect fine-grained temporal colour changes across short time intervals.

⁹ Zeki 1978.

¹⁰ Shapley and Hawkin 2011 713.

¹¹ Brewer, Liu, Wade and Wandell 2005.

¹² Shapley and Hawken 2011 714.

In any case, it is clear that at some point the human brain issues in categorical colour responses that are unaffected by changes in illumination. As I said right at the beginning, the precise range of such responses is not a matter to be decided behaviourally or introspectively, but will depend on facts of cognitive architecture. But we have seen no reason at all to posit anything like a million such possible responses, and reflection on their function argues that the range of such responses will be far more limited in number.

If I am right that there are two distinguishable cognitive mechanisms than make use of chromatic signals, then we should expect to find dissociations in injured subjects who possess one mechanism but not the other. And this is exactly what we do find.

Let me first consider the possibility of the difference detector without the surface classifier. Precisely this dissociation is arguably found in patients with the condition called ‘cerebral achromatopsia’. These are subjects who have suffered damage to the occipital ventral cortex near to the areas that have been conjectured to comprise the ‘colour centre’, and as a result have lost the ability to see hues.¹³ They report that everything is black and white and grey. Even so, they can easily detect boundaries between surfaces that differ only in transmitted wavelength and not in brightness.¹⁴ The natural explanation for this condition is that these subjects still possess the difference-detecting mechanism even though they have lost the surface classifier.

An interesting question which arises at this point is whether achromatopsics like this will still be capable of conscious judgements to the effect that one surface is more red, say, than another. The survival of some such ability might seem to be an implication of my argument so far. Even though the achromatopsics deny seeing hues, I am supposing that their chromatic difference detector is intact, and I have posited that this difference detector does not just issue in judgments that adjacent surfaces are chromatically different, but also pronounces on the relative location of the surfaces in various dimensions, including red-green and yellow-blue. The achromatopsics may not experience these differences as normal hue experiences of greater redness or yellowness, but this does not mean that they cannot display their sensitivity to them in other ways.

I do not know the answer to this question, nor whether there are any relevant empirical data¹⁵. I suspect it would not be easy to test for the relevant abilities. Many patients reported as suffering from cerebral achromatopsia have only partial loss of hue vision, and in addition many of them have various other cognitive deficiencies.¹⁶ In practice it could prove difficult to be sure which aspects of their performance on various tasks are due to which deficiencies.

In any case, there is no reason to be strongly committed to the specific suggestions I have made so far about the dimensions of comparative chromatic judgements issuing from the difference detector. For this mechanism to pronounce on which of two surfaces is redder, or yellower, or even brighter, there would need to be a clear division between double opponent cells in V1 that responded to the relevant channels of optical information, such as red-green, yellow-blue or bright-dark. However, it is not clear that the relevant double opponent cells are specialized in this way. Rather many of them seem to take in information from a mix of these channels (Shapley and Hawken 2011 706-8). If the V1 double-opponent cells do not display the relevant specialization, then the difference detector cannot be informing subjects about those specific dimensions of difference, and so must be making other qualitative comparisons, if it is making such comparisons at all.

¹³ Zeki 1990.

¹⁴ Heywood, Kentridge and Cowey 1998.

¹⁵ Achromatopsics standardly fail the Farnsworth-Munsell 100-hue test which requires the chromatic ordering of coloured chips (Heywood and Kentridge 2003 488). But this failure is arguably attributable to their loss of categorial hue responses, and does not necessarily show that they are incapable of any pairwise chromatic comparisons.

¹⁶ Davidoff 1997 119-22, Shapley and Hawken 2011 714.

Let me now turn to the possibility of the converse dissociation, the surface classifier without the difference detector. Given the difference detector's importance for form vision, subjects who lack it will be unable to discern three-dimensional form. (There is no serious possibility of subjects who have lost the chromatic difference detector still being able to discern form by means of brightness information alone. Even if, contrary the doubts expressed in the last paragraph, there are indeed double-opponent cells in V1 that are specialized for achromatic information, their spatial proximity to chromatic double opponent cells makes it likely that they will also be compromised by any lesions that incapacitate the latter.) The loss of form vision means that such subjects will be unable to perform colour constancy computations, since these rest on a prior identification of surfaces to which to assign uniform colours. Even so, a range of patients have been reported in whom the loss of form vision is combined with the retention of conscious hue experience.¹⁷ As we would expect, these patients cannot separate changes in background illumination from genuine changes in the colour of surfaces, and so experience the former as shifts in hue. The natural way to account for these patients is to suppose that their surface classifier is still operating, in that it is producing categorical colour responses when fed wavelength information, even in the absence of the fine-grained chromatic difference-detection needed to discern form.¹⁸

7. The Intransitivity of Indiscriminability. Further support for my alternative account of colour perception comes from reflection on the apparent intransitivity of colour indiscriminability. From the point of view of orthodoxy, this indiscriminability is extremely puzzling. My alternative account makes it much easier to understand.

In this section I shall explain the phenomenon and the puzzles it poses for orthodoxy. In the next section I shall show how my alternative point of view helps resolve these puzzles.

Consider a sequence of gradually varying colour samples where successive samples are below the threshold of a just-noticeable difference, but where the first and last are noticeably different. That is, sample₁ is indiscriminable from sample₂ when viewed side-by-side, and sample₂ similarly indiscriminable from sample₃, . . . and sample_{k-1} is similarly indiscriminable from sample_k, yet sample₁ and sample_k are discriminable when they are viewed side-by-side.

There is no doubt that it is possible to construct real examples of such sequences. However, if we assume the orthodox account of colour perception, they are very puzzling. On the orthodox account, a judgement that two samples viewed side-by-side are the same colour derives from the identity of the viewer's responses to each side: the awareness of sameness is a consequence, as we might put it, of 'the look of the left-hand side' being the same as 'the look of the right-hand side'. However, if this is what is going on, then it ought not to be possible to get sequences where the indiscriminability does not transmit from beginning to end. If the look of the first sample is the same as the look of the second, and the look of the second the same as that of the third, and so on, then can the look of the

¹⁷ Heywood and Kentridge 2003 494.

¹⁸ Certain cases of colour anomia, a condition where the ability to name colours is impaired, might also seem to indicate the presence of the difference detector without the surface classifier. In one thoroughly-studied such case, the patient retained hue vision, and was normal in many colour tasks, but performed oddly when asked to sort a range of colour chips from across the spectrum in 'whichever way he saw fit'. Where normal subjects would uncomplicatedly put the chips into around eight groups corresponding roughly to colour terms like red, yellow, brown, green, blue, purple and so on, this subject produced unusual groupings. His strategy was to pick up two chips and compare them to each other, placing them in the same group if they struck him as maximally similar. (Roberson, Davidoff and Braisby 1999.) The combination of these pairwise comparisons with the inability to form normal groupings might seem to call for explanation in terms of the difference detector without the surface classifier. However, given that this subject retained hue experience, and didn't report that everything looked black and white and grey, I am disinclined to suppose that he had lost the surface classifier, the outputs of which I am taking to constitute hue experience. Rather I would suggest that the subject's naming deficiency prevented him from grouping his conscious categorical colour responses into larger similarity classes in the way that normal subjects do.

first possibly be different from that of the last? Identity is a transitive relation, and this seems to leave no room for the looks of the first and last not to be the same.¹⁹

I shall be concerned in the rest of this paper with this very specific paradox about ‘colour looks’. It is worth distinguishing it from the more general phenomenon of sorites paradoxes. A sorites paradox arises whenever the first of a sequence of slightly differing objects $o_1, \dots, o_k$ satisfies some predicate ‘P’, and the last does not, yet there is a plausible tolerance principle of the form: if o_n is P, then o_{n+1} is P. In general, consistency in the face of such sorites sequences requires the rejection of the relevant tolerance principle; different theories of vagueness thus all agree in holding that certain instances of the relevant principles are not definitely true. But no corresponding move seems to be available with the ‘looks’ paradox generated by the orthodox view of colour discrimination. For this paradox does not depend on a tolerance principle, but on the simpler ‘identity assumption’ that:

- (I) sample_n and sample_{n+1} are visually indiscriminable if and only if the type of visual response produced by sample_n is the same as that produced by sample_{n+1}.

Note that the same sequences that generate our ‘looks’ paradox may well also generate a range of more familiar sorites paradoxes. Thus suppose that our sequence of samples takes us from a definitely red sample to a definitely orange one. Then ordinary sorites issues will be raised by the application of the predicates ‘red’ and ‘looks red’ to this sequence of samples, and also by the application of predicates like ‘experience as of red’ to the sequence of visual responses to the samples. I have nothing special to say about these sorites paradoxes. My concern, to repeat, is not with these paradoxes, but with the more specific ‘colour looks’ paradox generated by orthodox view of colour discrimination and the identity assumption it is committed to. (Note how the latter paradox will arise even if our sequence of samples remains within the definite oranges, or within any other nameable or conceptualizable colour category which might generate sorites problems, as long as the first sample is side-by-side distinguishable from the last. Our problem is not to do with transitions between categories, but with the way colour discrimination is assumed by orthodoxy to depend on identities and differences between ‘looks’.)²⁰

There are two main line of response to the ‘colour looks’ paradox in the literature, which I shall call the ‘standard line’ and the ‘minority line’.

The standard line accepts that the indiscriminability of colour samples is indeed intransitive, and resolves the paradox by modifying the identity principle (I).²¹ It continues to endorse one half of this principle—if two samples are discriminable, they must produce different conscious visual responses—but denies the other half, and holds that in some cases two samples can be indiscriminable even they do not produce the same conscious visual response. The idea here is that, as we move along the sequence, the conscious responses produced by at least one pair of adjacent samples must differ, in order to account for the difference in responses produced by the first and last samples. But this will be a difference in conscious responses that will not be reflected in the visual discriminability of the relevant pair of samples, since it is given that all adjacent pairs are visually indiscriminable.

The difficulty facing this line is obvious. In what sense can two samples have different conscious ‘looks’ if they are visually indiscriminable? Surely their indiscriminability means that they consciously look the same. The idea that these samples produce different conscious responses seems of dubious coherence. If the responses are consciously different, how can the subject be unaware of this?²²

¹⁹ A number of writers have argued on these grounds that the very idea of definite phenomenal properties is incoherent. Armstrong 1968, Dummett 1975.

²⁰ Rosanna Keefe (2011) emphasizes how specific solutions to the ‘looks paradox’ do not necessarily help with more general problems raised by vagueness.

²¹ Goodman 1951, Wright 1975, Williamson 2000

²² See Graff 2001 for a sustained argument that orthodoxy is committed to the transitivity of indiscriminability.

The minority line insists that genuine indiscriminability of conscious looks must indeed be transitive (after all, it is a matter of identity), and instead aims to deal with the paradox by denying that there are any genuine empirical sequences manifesting such intransitivity of indiscriminability²³. Adherents of this line allow that there are of course sequences of physical samples that are sequentially pairwise indistinguishable even though the first and last are distinguishable. But they deny that such sequences generate an intransitive sequence of indistinguishable ‘looks’, on the grounds that the same sample need not generate the same ‘look’ in different contexts. In particular, they suppose that at least one sample_k will generate one ‘look’ when compared with sample_{n-1}, but a different ‘look’ when compared with sample_{n+1}. This can then account for the difference in the ‘looks’ produced by sample₁ and sample_k, quite consistently with the transitivity of the identity of ‘looks’: since we aren’t guaranteed a sequence of identity of ‘looks’ as we go along the sequence of samples, there is no reason why the ‘look’ produced by the first and last sample should be the same.

Now, there are certainly many cases where the conscious colour response produced by some surface varies with context, as a wide range of contrast and adaptation effects testify. However, from the point of view of orthodoxy, it is not so obvious why this general phenomenon of context-dependence should apply here. The particular suggestion at hand is that some sample_n will shift appearance so as to look the same sample_{n-1} in one comparison, and the same as sample_{n+1} in another. But why should this happen? If sample_n is capable of looking different from sample_{n-1} in certain contexts (for example when it is compared with sample_{n+1}), then surely we would expect it to display this potential difference when it is actually juxtaposed to sample_{n-1}. Putting surfaces next to each other should surely make it easier to detect colour differences, not harder.

As we shall see in the next section, there is a sense in which I am happy to agree with the minority view about the effect of such juxtapositions on certain categorial classifications. However, my explanation of this phenomenon will hinge crucially on my non-orthodox view of the mechanisms of colour perception. Without such attention to mechanisms, it seems ad hoc to suggest that juxtaposition will make surfaces seem similar, rather than highlight their differences.

8. Intransitivity Resolved. The account of colour perception offered in this paper makes the apparent intransitivity of colour indiscriminability much easier to understand.

Note first that, on my account, the notion of ‘indiscriminability’ as applied to pairs of colour samples is ambiguous. On the one hand, it can mean that there is no ‘difference’ response when they are viewed side by side. On the other, it can mean that they are placed in the same colour category when viewed successively.

So let us distinguish between two samples being indiscriminable_{side-by-side} and their being indiscriminable_{singly}. Two samples are indiscriminable_{side-by-side} if the viewer can detect no difference when they are juxtaposed. Two samples are indiscriminable_{singly} if they produce the same colour classification when each viewed on their own.

First consider indiscriminability_{singly}. It seems clear that this must be transitive. After all, if a whole sequence of colour samples each produce the same categorial colour response when viewed singly, they so in particular must the first and the last.

But there is nothing at all problematic about this, since there is no reason to suppose that there are any empirical sequences that manifest intransitivity of indiscriminability_{singly}. Such a sequence would be one where sample₁ produced the same categorial response as sample₂ when they were viewed successively, and sample₂ similarly the same categorial response as sample₃, and so on . . . yet the last sample produced a different categorial response to the first sample. But there are not going to be any such sequences.

²³ Raffman 2000, Graff 2001.

To make it clear why, suppose that the first sample produces the categorial response orange and the last the categorial response yellow. Well, then there must be some point along the sequence where the viewer's categorial colour classifier switches from orange to yellow. It must switch somewhere, and that is where it does.²⁴ So it won't be the case that every pair of successors as we go along the sequence will be indiscriminable_{singly}. For some n , sample _{n} will produce one categorial colour response, but sample _{$n+1$} a different categorial response.

Now consider indiscriminability_{side-by-side}. Now there will certainly be empirical sequences of samples that manifest the intransitivity of indiscriminability_{side-by-side}—just manufacture a sequence of samples whose successive pairwise differences along some dimension of colour variation are below the threshold of 'just noticeability', but make the sequence long enough that the first and last will be noticeably different.

However, no paradox is generated by this intransitivity. We can simply accept that indiscriminability_{side-by-side} is indeed intransitive. That is simply how these side-by-side comparisons work—differences that are too small to see can add up, so to speak, to a visible difference.

This no longer generates any paradox because my gestalt view denies principle (I) and so does not allow the indiscriminable_{side-by-side} relation to impose a transitive equivalence relation on the samples.

The point is clearest in connection with an intransitive colour sequence all of whose members lie within some one colour category, say orange. The successive orange pairs are not noticeably different, but the first and last are. If the indiscriminability_{side-by-side} of the successive samples implied a sameness of 'looks', and the discriminability_{side-by-side} of the first and last a difference in 'looks', then we would face the apparent paradox discussed in the previous section. But on my gestalt view there is no such implication. Judgements of sameness_{side-by-side} and difference_{side-by-side} are neat gestalt judgements, not derived from comparisons of prior 'looks' responses to each sample. In particular, there is no inference from the difference_{side-by-side} of the first and last sample to their having different 'looks' when viewed singly. Indeed, these two samples, along with all the others, will actually share the only response evoked when they are viewed on their own, namely, the categorial response orange.

Still, what about the case where we move from one category into another via a series of indiscriminable_{side-by-side} comparisons—from from the oranges into the yellows, say?

Now it is given that the look_{singly} of the first sample is different from the look_{singly} of the last. The first is orange and the last is yellow. And this means, as above, that there must come a point, as we move along the sequence of samples, that some sample sample _{n} will be orange and the next sample sample _{$n+1$} will be yellow. Yet at the same time it is given that this same pair are indiscriminable_{side-by-side}—like all the successive pairs, their difference is below the threshold of just-noticeability.

This might seem very odd, but from my perspective it is still not paradoxical.

On the orthodox view, remember, judgements of discriminability_{side-by-side} are taken to derive from the difference or identity of the 'look' of each sample. And this leaves orthodoxy with no obvious way of explaining how the relevant samples can possibly be indiscriminable_{side-by-side} when the 'look' of one is orange and the 'look' of the next is yellow.

But from my perspective I can simply point out that the indiscriminability_{side-by-side} of these two samples, on the one hand, and their separate 'looks', on the other, are products of two separate mechanisms, the difference detector and the surface classifier respectively. The surface classifier will put the samples in two different boxes if they are presented separately (it has to switch somewhere

²⁴ The precise point at which the surface classifier switches may be affected by a number of factors, such as the order in which the samples are presented. This will not matter to any of the arguments that follow.

and this is where it does); while the difference detector will register the two samples as indiscriminable (because it is triggered by chromatic differences between adjacent samples and this difference is below its threshold).

This might be curious, but it is clearly not incoherent.

In one respect, my position is similar to the majority view that denies transitivity by rejecting half of the principle (I). That view, remember, upheld the intransitivity of indiscriminability by insisting that somewhere along the line the two samples in an indiscriminable pair must have different 'looks'. I concur, to the extent that I too hold there is a pair which is indiscriminable when viewed side-by-side but which will produce different categorial responses when viewed separately. The difference is that I do not see judgements of side-by-side indiscriminability as an upshot of the 'looks' produced when the samples are viewed separately. So I have an answer to the awkward question that so troubles the majority orthodox view, of explaining how the samples can possibly be indiscriminable, if they produce different 'looks'. My answer is that indiscriminability judgements depend on a different mechanism from the one that categorizes samples, and that this makes it perfectly possible for the former to detect no difference in some cases when the latter produces different classifications.

Here is an interesting question. How will it seem to you when you view the relevant pair of samples side-by-side? Will you both see the left-hand sample as orange and the right-hand sample as yellow, courtesy of the surface classifier, and see the two samples as indiscriminable in colour, courtesy of the difference detector? That would be strange, but we shouldn't rule it out a priori. There are other cases where distinct perceptual mechanisms produce judgements that contradict each other, for example the waterfall illusion (Crane 1988).

Still, I suspect that it won't work like this. I suspect that when the surface classifier is simultaneously presented with two side-by-side samples whose chromatic difference is below the just-noticeable threshold, it will always put them in the same colour category, even if these two samples mark the point where the surface classifier would switch categories if it were presented with a sequence of single samples successively. Think of it like this. Imagine now that you look at the whole sequence of side-by-side indistinguishable pairs of samples successively, starting with the first clearly orange pair, and ending with the last clearly yellow pair. Now again the surface classifier must switch from orange to yellow at some point. But if it is successively presented with the sequence of simultaneous pairs, then it seems likely that it will always switch between pairs, and never within them, not even within the pair of samples where it would have switched if the samples had been presented one-by-one.

If this is right, note that it follows that some particular sample_n will be put in the category orange when it is juxtaposed to sample_{n-1} and in the category yellow when it is juxtaposed to sample_{n+1}. So to this extent I am in a sort of agreement with the minority orthodox proposal; that the 'looks' of some samples changes with context. However my view, unlike that earlier proposal, can explain why this happens. The earlier puzzle, recall, was to explain why sample_n should be indistinguishable when juxtaposed with both sample_{n-1} and sample_{n+1}, if the brain is capable of seeing it as a different colour from one of them when not juxtaposed. If judgements of colour sameness/difference derive from independent judgements about each sample's individual colour, as orthodoxy has it, then this seems mysterious.

But if independent mechanisms are responsible for judgements of sameness/difference and judgements of categorial colour, as I have suggested, then the puzzle disappears. The reason sample_n is indistinguishable from both sample_{n-1} and sample_{n+1} is simply that the differences within each pair are too small to trigger the difference-detector. And the reason sample_n is classified as in different colour categories when juxtaposed with sample_{n-1} and sample_{n+1} respectively is simply that the surface-classifier isn't going to reach different verdicts on unnoticeably different samples when they are presented side-by-side.

9. Conclusion. The existence of two distinguishable mechanisms of colour perception, giving rise to categorial and relational colour responses respectively, is confirmed both by general functional considerations and by neurophysiological evidence. By distinguishing these two mechanisms, we can explain a number of puzzling features of colour perception, including the apparent ability of colour representation to outstrip conceptualization, and the intransitivity of colour indiscriminability.²⁵

References.

- Abramov, I. 1997 'Physiological mechanisms of color vision' in Hardin and Maffi 1997
- Armstrong, D. 1968 A Materialist Theory of the Mind London: Routledge and Kegan Paul
- Brewer, A., Liu, J., Wade, A., Wandell, B. 2005 'Visual field maps and stimulus selectivity in human ventral occipital cortex' Nature Neuroscience 8 1102-9
- Brewer, B. 1999 Perception and Reason Oxford: Oxford University Press
- Brewer, B. 2005 'Perceptual experience has conceptual content' in Sosa, E. and Steup, M. (eds) Contemporary Debates in Epistemology Oxford: Blackwell
- Chuard, P. 2006 'Demonstrative concepts without re-identification' Philosophical Studies 130 153-201
- Cohen, J. 2008 'Colour constancy as counterfactual' Australasian Journal of Philosophy 86 61 – 92
- Crane, T. 'The Waterfall Illusion' Analysis 48 142-7
- Dummett, M. 1975 'Wang's Paradox' Synthese 30 201-32
- Evans, G. 1982 The Varieties of Reference Oxford: Oxford University Press
- Goodman, N. 1951 The Structure of Appearance Cambridge, Mass: Harvard University Press
- Graff, D. 2001 'Phenomenal continua and the sorites' Mind 110 905-35
- Hamwi, V. and Landis, C. 1955 'Memory for colour' Journal of Psychology 39 183-94
- Hardin, C. And Maffi, L. (eds) Colour Categories in Thought and Language Cambridge: Cambridge University Press 1997
- Heywood, C. and Kentridge, R. 2003 'Achromatopsia, color vision & cortex' Neurological Clinics of North America 21: 483-500

²⁵ I have given versions of this paper at King's College London, the Institute of Philosophy, Edinburgh, Warwick, Stockholm, Queen's University Belfast, Cambridge, Rutgers, and the Phenomenal Qualities Project at Hertfordshire. I would like to thank all those who made comments on those occasions. I can particularly remember being helped by Peter Adamson, Charles Travis, Paul Coates, Philip Goff, Sam Coleman, Guy Longworth, Kranti Saran, Naomi Eilan, Åsa Wikforss, Kathrin Glüer-Pagin, Peter Pagin, Sören Häggqvist, Levi Spectre, Tim Crane, Will Davies, Hugh Mellor, Barry Loewer, Brian McLaughlin, Barry Smith, David Chalmers, Michael Tye, Mike Martin, Ophelia Deroy and Fiona Macpherson. I have also benefitted from exchanges with Jonathan Cohen, Ned Block, Pete Mandik, Kevin O'Regan, Jonathan Salisbury, Mohan Matthen, Yasmina Jraissati and Diana Raffman (whose ingenious arguments in her 1995 paper originally provoked my view of colour perception).

- Heywood, C., Kentridge, R. and Cowey, A. 1998 'Cortical color blindness is not 'blindsight for color'' Consciousness and Cognition 7 410-42
- Hurvich, L. and Jameson, D. 1957 'An opponent-process theory of color vision' Psychological Review 64 384-404
- Keefe, R. 2011 'Phenomenal sorites paradoxes and looking the same' Dialectica 65 327-44
- Kelly, S. 2001 'Demonstrative concepts and experience' The Philosophical Review 110 397-420.
- Kentridge, R., Heywood, C. & Davidoff, J. 2003 'Color perception' in Arbib, M. (ed) Handbook of brain theory & neural-networks (2nd Edition) Cambridge Mass: MIT Press 230-233
- McCamy, C. 1998 'The number of discernible colours' Colour Research and Application 23 337
- McDowell 1994. Mind and World Cambridge Mass: Harvard University Press
- Mandik, P. 2012 'Color-conscious conceptualism' Consciousness and Cognition 21 617-31
- Peacocke, C. 1992. A Study of Concepts Cambridge Mass: MIT Press.
- Pointer, M. and Attridge, G. 1998 'The number of discernible colours' Colour Research and Application 23 52-54
- Raffman, D. 1995 'On the Persistence of Phenomenology' in Metzinger, T. Conscious Experience Paderborn: Schoenigh Verlag
- Raffman, D. 2000 'Is perceptual indistinguishability nontransitive?' Philosophical Topics 28 153-74
- Regier, T. and Kay, P. (2009). 'Language, thought, and color: Whorf was half right' Trends in Cognitive Sciences 13 439-46
- Roberson, D., Davidoff, J. and Braisby, N. 1999 'Similarity and categorisation: neuropsychological evidence for a dissociation in explicit categorisation tasks' Cognition, Vol 71 1-42
- Roberson J., Davidoff, J., Davies, I. and Shapiro, L.. 2004. 'The development of color categories in two languages: a longitudinal study' Journal of Experimental Psychology: General 133 554-71
- Roberson, D., Davidoff, J. Davies, I. and Shapiro, L. 2005. "Color categories: evidence for the cultural relativity hypothesis" Cognitive Psychology 50 378-411.
- Rosenthal, D. 2005 Consciousness and Mind Oxford: Oxford University Press
- Shapley, R. and Hawken, M. 2011 'Color in the cortex: single- and double-opponent cells' Vision Research 51 701-17
- Sivic, L. 'Color Systems for Cognitive Research' 163-93 in Hardin and Maffi 1997
- Werner, J. S., and Wooten, B. R. 1979 'Opponent chromatic mechanisms: relation to photopigments and hue naming' Journal of the Optical Society of America 69 422-434.
- Williamson, T. Knowledge and its Limits Oxford: Oxford University Press

Winawer, J., Witthoft, N., Frank, M. ,Wu, L., Wade, A., and Boroditsky, L. 2007 'The Russian blues reveal effects of language on color discrimination' Proceedings of the National Academy of Science 104 7780-5

Wright, C. 1975 'On the Coherence of Vague Predicates' Synthese 30 325-65

Zeki. S. 1978 'Functional specialization in the visual cortex of the rhesus monkey' Nature 274 423-8

Zeki, S. 'A century of cerebral achromatopsia' Brain 113 1721-77