

Colour constancy and Fregean representationalism

Boyd Millar

Abstract All representationalists maintain that there is a necessary connection between an experience's phenomenal character and intentional content; but there is a disagreement amongst representationalists regarding the nature of those intentional contents that are necessarily connected to phenomenal character. Russellian representationalists maintain that the relevant contents are composed of objects and/or properties, while Fregean representationalists maintain that the relevant contents are composed of modes of presentation of objects and properties. According to Fregean representationalists such as David Chalmers and Brad Thompson, the Fregean variety of the view is preferable to the Russellian variety because the former can accommodate purported counterexamples involving spectrum inversion without illusion and colour constancy while the latter cannot. I maintain that colour constancy poses a special problem for the Fregean theory in that the features of the theory that enable it handle spectrum inversion without illusion cannot be extended to handle colour constancy. I consider the two most plausible proposals regarding how the Fregean view might be developed in order to handle colour constancy—one of which has recently been defended by Thompson (Australas J Philos 87:99–117, 2009)—and argue that neither is adequate. I conclude that Fregean representationalism is no more able to accommodate colour constancy than is Russellian representationalism and, as such, ought to be rejected.

Keywords Representationalism · Fregean representationalism · Russellian representationalism · Phenomenal consciousness · Perceptual content

Representationalism is the view that an experience's intentional content (what the experience is about or what it represents) and its phenomenal character (what the experience is like for the subject) are necessarily connected: any two experiences that have a certain intentional content have a certain phenomenal character and any two experiences that have that phenomenal character also have that intentional content. However, amongst representationalists there is a dispute regarding the nature of the contents that are necessarily connected to phenomenal character. Russellian representationalists maintain that the relevant contents are composed of objects and/or properties, while Fregean representationalists maintain that the relevant contents are composed of *modes of presentation* of objects and properties.¹

An important potential reason to prefer the Fregean view to the Russellian view is that the Fregean view has resources for responding to purported counterexamples to representationalism that the Russellian view lacks. One standard method of demonstrating that representationalism is false is to point to two experiences that are phenomenally identical but possess different intentional content. There are two general types of example that purport to do just this. First, there are inverted spectrum and twin earth scenarios that generally involve two different subjects and/or two different environments (I'll call these Type A examples).² Second, there are real-world examples, such as those involving the phenomenon of colour constancy, that are restricted to a single subject in a single environment (I'll call these Type B examples).³ Fregean representationalists such as Chalmers (2004, 2006) and Thompson (2008, 2009) have argued convincingly that Fregean representationalism is preferable to Russellian representationalism because the former can accommodate Type A examples while the latter cannot. Moreover, Thompson (2006, 2009) has argued that Fregean representationalism has the added advantage that it can accommodate Type B examples while Russellian representationalism cannot.

I maintain, to the contrary, that Type B examples constitute a special problem for Fregean representationalism. My focus here will be examples involving colour constancy: I'll argue that while the Fregean view has an advantage over the Russellian view in that it can accommodate Type A examples, it is no more able to accommodate examples of colour constancy. First, in Sect. 1 I will say a bit more about the difference between Fregean and Russellian representationalism and explain how the Fregean view is able to accommodate Type A examples such as spectrum inversion without illusion. Next, in Sect. 2 I will explain why colour constancy presents a particular challenge for Fregean representationalism. Then, in Sects. 3 and 4 I will consider what I take to be the two most plausible proposals regarding how the Fregean view might be developed in order to handle colour constancy—the first of which was recently defended by Thompson (2009), and the second is closely related to his view—and argue that neither is adequate. I conclude that Fregean representationalism is no more able to accommodate colour constancy than is Russellian representationalism and, as such, ought to be rejected.

¹ These two views will be explained in further detail in Sect. 1 below.

² See, for example, Block (1990).

³ A well known example of this sort is Peacocke's (1983, pp. 16–17) three-dimensional Necker cube.

1 Spectrum inversion without illusion

In order to explain the difference between Fregean and Russellian representationalism it will help to focus on phenomenally red visual experiences. When you look at a ripe tomato, a Chinese flag, a red wall under white light, or a white wall under red light, these various visual experiences are phenomenally similar in a certain obvious respect; that common aspect of the phenomenal character of these different experiences is a particular phenomenal property we can call *phenomenal redness*. A representationalist maintains that every experience that instantiates phenomenal redness possesses a certain specific intentional content (and every experience that possesses that content instantiates phenomenal redness); however, Fregean and Russellian representationalists disagree regarding the nature of the intentional content at issue.

A Russellian representationalist maintains that every visual experience instantiating phenomenal redness possesses an intentional content composed of a certain property.⁴ More specifically, the standard Russellian view is that every phenomenally red experience represents that the relevant object in the environment possesses a certain physical property, such as a specific spectral reflectance.⁵ The Fregean representationalist, on the other hand, maintains that every visual experience instantiating phenomenal redness has a certain *mode of presentation* as its content. As Chalmers (2004, p. 172) explains it, modes of presentation should be thought of here as “conditions on extension”: conditions that a property must satisfy in order to be the property represented by the relevant experience. With respect to the present case, Fregean representationalists maintain that every phenomenally red experience has as its content a mode of presentation that picks out the property that normally causes phenomenally red experiences in the subject.⁶

This account of the nature of the intentional contents that are necessarily connected to specific phenomenal properties gives the Fregean representationalist an important advantage over the Russellian representationalist when it comes to accommodating Type A examples. For instance, in a standard inverted spectrum scenario there are two subjects, call them Jack and Jill, whose colour phenomenology is red-green inverted relative to one another. Suppose that when Jack looks at a red object and Jill looks at a green object each has a phenomenally red colour experience; nonetheless, we have a strong intuition that neither subject pervasively misperceives the colours of objects. According to Russellian representationalism, however, any two phenomenally identical colour experiences must attribute the same physical colour properties to perceived objects. Since, then, the present scenario involves different coloured objects, the Russellian is forced to say that at least one of the subjects’ experiences misrepresents its respective object.

⁴ See Chalmers (2004, p. 167).

⁵ Chalmers (2004, p. 168) calls this specific view “physical Russellian representationalism”; Thompson (2008, p. 204) calls it “Standard Russellianism.” In what follows, I will be concerned with this standard Russellian view exclusively, so from now on whenever I say “Russellian representationalism” I mean the standard variety of Russellian representationalism. For some of the difficulties facing non-standard varieties of Russellian representationalism see Chalmers (2004, pp. 168–171) and Thompson (2007).

⁶ See Chalmers (2004, p. 176) and Thompson (2009, p. 102).

The Fregean representationalist can avoid this difficulty because Jack's and Jill's phenomenally identical experiences don't have to attribute the same colour properties to their respective objects—they only have to share a mode of presentation. The Fregean can say that when Jack looks at a red object and Jill looks at a green object each has a colour experience with the same mode of presentation as its content: a mode of presentation like "*the property that normally causes phenomenally red experiences in me*" (Chalmers 2004, p. 176). But when Jack has an experience representing that a certain object has *the property that normally causes phenomenally red experiences in me*, that experience attributes the property of physical redness to the relevant object; whereas when Jill has an experience with this same content, that experience attributes the property of physical greenness to the relevant object, since physical greenness is the property that typically causes phenomenally red experiences in her. Consequently, even though when Jack looks at a red object and Jill looks at a green object each has a phenomenally red colour experience, both of these experiences can be veridical representations of the relevant object's colour.⁷

2 The problem of colour constancy

While Fregean representationalism thus possesses a significant advantage over Russellian representationalism in that it accommodates spectrum inversion without illusion quite naturally, similar examples restricted to a single subject in a single environment still pose a special problem for the view. A simple way to generate examples of this sort is to appeal to the fact that objects look to have a single, constant colour under a variety of different lighting conditions—a psychological phenomenon known as colour constancy. Due to colour constancy, experiences of different objects under different lighting conditions can have the same phenomenal character and yet represent the relevant objects to be different colours. Thompson (2009, p. 110) describes a particular example of this sort. First, consider your experience as you view a white sheet of paper in the yellowish sunlight of the early morning. Next, consider your experience as you view a slightly yellowish sheet of paper under fluorescent light. If the conditions are just right, these two experiences of the two different sheets of paper will both have the very same phenomenal character: both will have a certain determinate colour phenomenology which, for the sake of simplicity, I will call "phenomenal yellowness."⁸ Yet, there is no reason

⁷ The same feature enables Fregean representationalism to accommodate inverted earth scenarios: see Thompson (2009, pp. 104–105).

⁸ One might be tempted to resist the claim that these two experiences of the sheets of paper have the same colour phenomenology, but there doesn't seem to be any aspect of colour phenomenology that they don't share. As Thompson argues, "the phenomenal colour appearance of a place in the visual field is exhaustively characterized by hue, saturation, and brightness (understood as features of visual experience, not as physical qualities of light). This is not only introspectively verifiable—it is the standard view in the literature on colour perception" (2006, p. 84). And he notes further that in colour constancy experiments [for instance, Arend, Reeves, Schirillo and Goldstein (1991)], subjects are able to match two coloured papers for hue, saturation, and brightness even when those same two papers do not look to the subject to have the same surface colour (Thompson 2006, p. 85).

to assume that in either case your experience will misrepresent the colour of the paper—in fact, we know that the visual systems of normal humans are very good at determining the unchanging colour of a surface despite changes to its “sensory appearance.”⁹ So, once again, we have two phenomenally identical experiences of different coloured objects and yet each experience represents the colour of the relevant object veridically.

This example constitutes a problem for Russellian representationalism for precisely the same reasons that the inverted spectrum scenario does.¹⁰ However, the feature of Fregean representationalism that enables it to accommodate spectrum inversion without illusion—building the *typical causal basis* of the relevant phenomenology in the subject of the experience into the mode of presentation—cannot be extended to handle colour constancy. What typically causes phenomenally yellow experiences in you might be looking at objects that are a certain specific shade of yellow; so if the relevant mode of presentation picked out the property that typically causes phenomenally yellow experiences in you, then your experience of the white sheet of paper in yellow sunlight would misrepresent the paper to be that shade of yellow.

Consequently, the Fregean representationalist needs to add something to his account of the relevant modes of presentation. The most natural suggestion would be to build in lighting conditions; but that suggestion turns out to be unworkable. First, you can't include the specific lighting conditions in a particular situation. For instance, when you view the white sheet of paper under yellow light the mode of presentation can't pick out the property that typically causes phenomenally yellow experiences under yellow light; if that were the correct characterization of the content of your experience then your experience would of course be veridical, but that content could not be shared by your phenomenally identical experience of the yellow sheet of paper under fluorescent light. To overcome this difficulty one might suggest that the mode of presentation specifies lighting conditions indexically—that is, the relevant mode of presentation might pick out the property that typically causes phenomenally yellow experiences in the subject under *present* lighting conditions. However, as Thompson explains (2009, pp. 110–111), such a suggestion is unacceptable because it would make misperceiving an object's colour virtually impossible.¹¹

3 Thompson's solution

So, the Fregean representationalist has a problem: she needs to add something to the relevant modes of presentation in order to accommodate colour constancy, but the most obvious strategy (including lighting conditions) won't work. A more

⁹ See, for example, Arend et al. (1991).

¹⁰ See Thompson (2006).

¹¹ Another possibility would be to claim that the mode of presentation picks out the property that typically causes phenomenally yellow experiences in the subject under those lighting conditions *represented* to obtain by the present experience. Thompson (2009, pp. 111–112) provides convincing reasons for rejecting this proposal as well.

promising way for the Fregean to respond to this difficulty, recently defended by Thompson, would be to suggest that “phenomenal modes of presentation for colour experiences might be sensitive to properties of the entire visual field in their determination of reference” (2009, p. 112). In other words, the proposal is that the mode of presentation for a phenomenally yellow experience picks out the property that typically causes phenomenally yellow experiences in the subject in instances where the total visual phenomenology of the experience is similar to the total visual phenomenology of the present experience.

For instance, when you view the white paper in the early morning the yellowness of the light will be evident from the phenomenology of other areas of your visual field—the things that happen to be surrounding the sheet of paper will have the yellowish appearance typical of things illuminated by yellow light. When you have experiences the total visual phenomenology of which is similar to this experience, phenomenal yellowness at a particular location in the visual field is typically caused by a white object, and so your experience of the paper in this case represents that it is white rather than yellow. On the other hand, while when considered in isolation the slightly yellowish sheet of paper illuminated by fluorescent light has the same phenomenal character as the white paper illuminated by yellow light, the total visual phenomenology of this experience will be different. Specifically, the surrounding areas of your visual field will have the appearance typical of things illuminated by fluorescent light, and in the context of such phenomenology the phenomenally yellow experience of the paper will represent that it is a certain specific shade of yellow.

Thompson’s proposal, then, allows for the experience of the white paper in yellow light and the experience of the yellow paper in fluorescent light to attribute different colour properties to their respective objects since there will be differences in the total visual phenomenology of the two experiences. In addition, by making modes of presentation sensitive to the phenomenology of the entire visual field, his proposal avoids the objection that it makes misperceiving an object’s colour impossible. For instance, Thompson (2009, p. 110) describes a case where there are two objects on a stage, a red object illuminated by normal incandescent light and a white object illuminated by a carefully positioned red spotlight. Assuming that conditions are controlled such that the redness of the spotlight leaves “no phenomenal trace in other regions of the visual field,” an experience of this scene would misrepresent the white object to be red, since when you have experiences the total visual phenomenology of which is similar to this experience, phenomenal redness at a particular location in the visual field is typically caused by a red object.

While I think that Thompson’s proposal constitutes the best hope the Fregean theory has of accommodating colour constancy, I believe it is ultimately unacceptable. He assumes that any difference in lighting conditions that makes a difference to the colour properties your experience attributes to objects will be accompanied by some difference in the total visual phenomenology of your experience; but this assumption is not obviously correct. The difficulty, then, is that his proposal has the consequence that “phenomenally identical total visual experiences within a single subject” cannot represent objects to be different colours (Thompson 2009, p. 113). But there would seem to be counterexamples to this claim.

For instance, imagine a case where you're in a room with white walls illuminated from above by a diffuse yellow light. Imagine too that there are a variety of different coloured objects in this room and that you spend a good deal of time comparing the colour of the walls to the colour of these various objects. In such a situation, just as in Thompson's original case of the white piece of paper viewed in yellow light, when you look at a given wall your visual experience will instantiate phenomenal yellowness and yet accurately represent that the wall is white. Now imagine that you walk towards one of the walls in this room, getting close enough so that it encompasses your entire visual field. Presumably in such a circumstance the colour of the wall will not look to have changed: as you get close enough to the wall such that it fills your entire visual field, you will not be subjected to a visual illusion to the effect that the colour of the wall has changed from white to yellow.¹² Consequently, even though the same phenomenal yellowness now fills your entire visual field, your visual experience will accurately represent that the wall is white.

Next, imagine a case where someone with sufficient knowledge of colour and light repaints the walls in this room a certain shade of yellow and replaces the yellow light with white light, such that when you look at the walls again your visual experience instantiates the very same phenomenal yellowness. Imagine too that once again there are a variety of different coloured objects in this room and that you spend a good deal of time comparing the colour of the walls to the colour of these various objects. In such a situation, just as in Thompson's case of the yellow piece of paper viewed under fluorescent light, when you look at a given wall your visual experience will instantiate phenomenal yellowness and accurately represent that the wall is yellow. Next, imagine that you again walk towards one of the walls in the room, getting close enough that it encompasses your entire visual field. In such a circumstance the colour of the wall will not look to have changed: as you get close enough to the wall such that it fills your entire visual field your visual experience will accurately represent that the wall is yellow.

These two experiences of the different coloured walls that fill your entire visual field thus constitute a counterexample to Thompson's proposal. Both experiences instantiate the same phenomenal yellowness throughout the entire visual field; but since the colour of the wall is different in the two circumstances, and since we have good reasons to think that the colour of the wall will be represented accurately in each case, your two experiences will represent different colours despite possessing precisely the same total visual phenomenology.

One way for the Fregean representationalist to respond would be to suggest that these two experiences can possess the same content and both be accurate (despite

¹² If you hadn't first been able to look around the room at a variety of different coloured objects, there would be no way for your visual system to distinguish the colour of the wall from the colour of the light [see, for example, Kaiser and Boynton (1996, pp. 510–511)]; however, when the nature of the light is not apparent in a given experience, the visual system can rely on information from immediately prior experiences to achieve colour constancy [see, for example, Smithson and Zaidi (2004)]. It has also been shown that even in cases where only a single surface is visible a good deal of colour constancy can be achieved so long as the subject has been exposed to the relevant illuminant for a sufficient length of time (Uchikawa et al. 1989).

the change to the actual colour of the wall) so long as the content in question is sufficiently vague.¹³ In other words, the suggestion would be that the two experiences in the present example both represent the colour of the wall to fall within a specific range of colours and that each time the actual colour of the wall falls within that range. However, I think that such a suggestion is clearly incompatible with the phenomenology of colour experience: the surfaces of objects look to have highly determinate colours, they don't look to have a colour falling anywhere within a certain range of colours. In ordinary circumstances, if someone asked you to describe the colour of a wall you would never say "the colour looks to be anywhere between a pure white and a rather warm off-white." One might reply that in the case at issue circumstances are not ordinary because the wall fills your entire visual field. However, it is simply false that your experience represents colour with less specificity when there are fewer colours visible to you. To test this for yourself, try walking towards a wall until it fills your entire visual field. Did the wall look to have a highly specific shade when you could see a number of different objects of different colours? And when you were close enough that the wall filled your entire visual field, did it no longer look to be that same specific shade—did it instead look as though the colour merely fell somewhere within a certain range of shades? Hopefully you will agree that the colour of the wall was not represented with any less specificity once it filled your entire visual field.

Assuming, then, that both experiences in the present scenario represent the colour of the wall accurately and that both represent the colour to be a specific shade, we have an example where two experiences have the same total visual phenomenology and yet attribute different colours to the surface of a given object. Since Thompson's proposal has the consequence that such cases are impossible, we ought to reject Thompson's proposal.

4 The remaining alternative

Essentially, Thompson's suggestion is that what specific colour property an experience with a certain colour phenomenology attributes to an object depends on the context of the total visual phenomenology of the experience at the time in question (and so he builds this context into the relevant modes of presentation). A related proposal that might escape the objection just described would be to suggest that not only does the specific colour property an experience with a certain colour phenomenology attributes to an object depend on the context of the total visual phenomenology of the experience at the time in question, but it also depends on the total visual phenomenology of *prior experiences*.¹⁴ The proposal would be that, for example, the mode of presentation for a phenomenally yellow experience picks out the property that typically causes phenomenally yellow experiences in the

¹³ This possible response was suggested to me by Brad Thompson.

¹⁴ This proposal was suggested to me by Brad Thompson. Jagnow (2009, 566 n) also notes that a Fregean might need to adopt such a view in order to address the difficulty at issue.

subject in instances where the total visual phenomenology of the experience is similar to the total visual phenomenology of the present experience and where the total visual phenomenology of the experiences preceding that experience for some period of time is similar to the total visual phenomenology of those preceding the present experience.

A Fregean representationalist who characterizes the relevant modes of presentation along these lines might be in a position to accommodate the example described in the previous section since the colour that your experience represents the wall to be each time it fills your entire visual field clearly depends on the phenomenology of your preceding visual experiences. However, such a proposal would face the difficulty that the nature of the light in a given circumstance might be made apparent to you in any number of different ways. Accordingly, there could be a circumstance where the specific sequence of experiences leading up to a particular experience might be totally novel, and yet that experience might represent the colour of the relevant object accurately. Such a case would be incompatible with the present suggestion because there would be no property that satisfied the proposed condition: according to the present proposal, any visual experience preceded by a novel sequence of experiences would lack a truth value.¹⁵

But even leaving this objection to one side, I think there are straightforward counterexamples to the present Fregean proposal. This proposal has the consequence that two experiences that are alike both in their present total visual phenomenology and the total visual phenomenology of the experiences for some period of time preceding them cannot represent objects to be different colours. But, this consequence would seem to be false.

Consider the fact that a given stimulus can be ambiguous as to whether some part of the stimulus is a darker colour than the rest, or whether some part of the stimulus is in shadow.¹⁶ Consequently, there might be cases where you are presented with such a stimulus and your visual system has to choose between the two “interpretations,” just as it does when you look at an ambiguous figure like the Necker cube or the duck/rabbit figure. You can create a stimulus of this sort if you take a cube each face of which is the same colour, position it so that you can see only the front and top faces, and light it from the front such that there is a shadow attached to the top face. If you carefully control the light such that there are no visual cues as to the direction of the light beyond the darker appearance of the top face of the cube (for example, no other shadows should be visible from your vantage point), when you look at the cube you should be able to switch back and forth between seeing the top face as darker than the front face and seeing the top face as the same colour as the front face but in shadow.

¹⁵ Thompson (2009, pp. 105–106) points out that his view has precisely this consequence in the case of a subject’s first experience of any given colour. According to Thompson, such experiences possess contents but their contents lack truth values: “They have conditions of satisfaction... But those conditions of satisfaction involve a presupposition that is not satisfied—that there *is* a typical cause of experiences of that phenomenal type. This makes them infelicitous, but not false” (2009, p. 106).

¹⁶ This ambiguity is noted regularly by researchers: see, for example, Gilchrist (1979, p. 115), Cavanagh and Leclerc (1989, p. 4) and Ostrovsky et al. (2005, p. 1301). And such ambiguous stimuli are sometimes used in experiments: see Jacobson and Werner (2004).

However, there will be no phenomenal difference anywhere in the visual field between these two experiences.¹⁷

For present purposes, the crucial point is that the nature of your visual phenomenology prior to each experience *need not* influence the “interpretation” your visual system settles on. For instance, imagine a scenario where you’re part of a carefully controlled experiment where you’re kept in the dark for a long period of time right up until the moment you’re presented with the ambiguous stimulus. Nonetheless, if you were subjected to the experiment on two different occasions it would still be possible for your visual experience to represent the scene differently each time: the first time you’re presented with the cube it might look as though the top face is darker than the front face; and the second time it might look as though both faces are the same colour but the top face is in shadow (after all, you can get a similar outcome if you look at the Necker cube on two separate occasions).

Assuming there is no phenomenal difference anywhere in the visual field between these two experiences and assuming that the phenomenal character of the experiences preceding each could be exactly the same, the only other way for a defender of the present Fregean proposal to respond to the example would be to insist that the two experiences attribute the same colours to the same parts of the cube. Someone responding this way will have to say that there is no difference with regard to the intentional content of your two visual experiences of the ambiguous stimulus—the only thing that changes is that you make a different *judgment* on the basis of your visual experience each time.

Someone who made this reply would have to say one of two things: either both visual experiences of the cube represent the top face to be a darker colour than the front face and you merely judge on one occasion that both faces are the same colour and the top face is in shadow; or both visual experiences of the cube represent that both faces are the same colour and the top face is in shadow and you merely judge on one occasion that the top face is darker than the front face. However, neither account of the present example is plausible. First, given what we know about the neural mechanisms underlying colour constancy and given the evidence of colour constancy in non-human animals, it’s implausible to suggest that distinguishing darker colours from shadows requires that one form judgments on the basis of one’s visual experience.¹⁸

Second, according to each of these two accounts, one of the times you’re presented with the ambiguous stimulus you form a judgment that’s inconsistent with the content of your visual experience. Of course, we sometimes form such judgments; but it’s problematic to appeal to such judgments in response to the present example. First, if one claims that both times you look at the cube your visual experience represents that both faces are the same colour and the top face is in shadow and that you merely judge on one occasion that the top face is darker than the front face, then on that occasion you form a judgment that is inconsistent with

¹⁷ When you view a standard ambiguous figure such as the Necker cube it’s arguable that the visual phenomenology of your experience changes as you switch from one “interpretation” to the other. But in such cases there is a dramatic reorganization of the visual field that is completely absent in the present example.

¹⁸ This point is made forcefully by Thompson (2009, p. 109).

the content of your visual experience. The difficulty here is that people do not form judgments inconsistent with the content of their visual experiences unless they have independent reasons for doing so (in other words, you need a reason for not believing your eyes). For example, if you were completely unfamiliar with the Müller-Lyer diagram, when you first viewed it you would certainly judge that the two lines were different lengths. It is not until you measure the lines, or until someone tells you it is a trick, that you judge the lines to be the same length. However, with respect to the present example there is no necessity to presuppose that you have any independent reasons for judging, contrary to the content of your visual experience, that the top face of the cube is a darker colour than the front face: even if you have no such reasons you can switch between the two different “interpretations.”

Alternatively, if one claims that both times you look at the cube your visual experience represents that the top face of the cube is a darker colour than the front face but on one occasion you judge that both faces are the same colour, then on that occasion you form a judgment that is inconsistent with the content of your visual experience. The reason this claim is implausible is that typically when you judge contrary to the content of your visual experience, you are immediately aware of the conflict between what your eyes are telling you and what you judge to be the case. Once again, consider what happens when you fall victim to the Müller-Lyer illusion. When you measure the two lines and *judge* them to be equal in length, one line continues to *look* longer than the other; but in such a case you are immediately aware that the content of your visual experience is inconsistent with the content of your judgement. According to the present account of our cube example, then, when you judge that the top face of the cube is the same colour as the front face but in shadow, you ought to be aware of a conflict between that judgment and the content of your visual experience; but, of course, you are aware of no such conflict.

Finally, to insist that it's only the judgments one forms that change when one looks at the cube is to deny that it's possible to create a stimulus that is ambiguous in the manner described (after all, it would be clearly false to claim that when you look at an ambiguous figure like the Necker cube the content of your visual experience is always the same, but you make different judgments on different occasions). But, given that we know that different visual experiences can attribute different spatial properties to the same stimulus viewed under the same conditions (as the Necker cube shows), it would be rather implausible to insist that different visual experiences can never attribute different colour properties to the same stimulus viewed under the same conditions.

5 Conclusion

Ultimately, then, it appears that colour constancy does indeed constitute a special problem for Fregean representationalism. The Fregean representationalist accommodates colour inversion without illusion by maintaining that the content that's necessarily shared by any experience with a specific colour phenomenology is a mode of presentation that picks out the property that typically causes experiences

with that phenomenology in the subject of the experience. But, thanks to colour constancy, a subject can have a veridical colour experience even when (due to atypical lighting conditions) the colour in question is not the property that typically causes experiences with the relevant phenomenology in the subject of the experience. In order to accommodate colour constancy, then, the characterization of the relevant modes of presentation must be modified in some way. But, as we've now seen, the most plausible proposals—including lighting conditions, including the phenomenology of other regions of the visual field, or including the phenomenology of prior experiences—are all unacceptable. As such, it's reasonable to conclude that Fregean representationalism simply cannot accommodate colour constancy and therefore ought to be rejected.

Perhaps some form of representationalism yet more complex than those considered above could accommodate colour constancy. But the mere fact that an ordinary, everyday perceptual phenomenon like colour constancy constitutes such a significant challenge for representationalism should make us suspicious of the view. It may be that continuing to search for some modification to representationalism that makes the view consistent with colour constancy is not worth the effort. Examples of colour constancy show that two experiences with the same colour phenomenology can nonetheless represent their respective objects to have different colours (and two experiences with different colour phenomenology can represent their respective objects to have the same colour). And it seems quite natural to assume that two colour experiences that represent their respective objects to have different colours differ with regard to their intentional content. Thus, colour constancy seems to provide a powerful motivation to deny representationalism altogether and explore the view that a perceptual experience's phenomenology can vary independently of its intentional content.¹⁹

Acknowledgment I am grateful to Brad Thompson for comments on an earlier version of this material.

References

- Arend, L., Reeves, A., Schirillo, J., & Goldstein, R. (1991). Simultaneous color constancy: Papers with diverse Munsell values. *Journal of the Optical Society of America A*, 8, 661–672.
- Block, N. (1990). Inverted Earth. In J. Tomberlin (Ed.), *Philosophical perspectives, 4: Action theory and philosophy of mind* (pp. 53–79). Atascadero, CA: Ridgeview.
- Cavanagh, P., & Leclerc, Y. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 3–27.
- Chalmers, D. (2004). The representational character of experience. In B. Leiter (Ed.), *The future for philosophy* (pp. 153–181). Oxford: Clarendon Press.
- Chalmers, D. (2006). Perception and the fall from Eden. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 49–125). Oxford: Clarendon Press.
- Gilchrist, A. (1979). The perception of surface blacks and whites. *Scientific American*, 240, 112–124.
- Jacobson, J., & Werner, S. (2004). Why cast shadows are expendable: Insensitivity of human observers and the inherent ambiguity of cast shadows in Pictorial Art. *Perception*, 33, 1369–1383.
- Jagnow, R. (2009). How representationalism can account for the phenomenal significance of illumination. *Phenomenology and the Cognitive Sciences*, 8, 551–572.

¹⁹ I defend a version of this view in Millar (2011).

- Kaiser, P., & Boynton, R. (1996). *Human color vision* (2nd ed.). Washington, DC: Optical Society of America.
- Millar, B. (2011). Sensory phenomenology and perceptual content. *Philosophical Quarterly*, 61, 558–576.
- Ostrovsky, Y., Cavanagh, P., & Sinha, P. (2005). Perceiving illumination inconsistencies in scenes. *Perception*, 34, 1301–1314.
- Peacocke, C. (1983). *Sense and content*. Oxford: Clarendon Press.
- Smithson, H., & Zaidi, Q. (2004). Colour constancy in context: Roles for local adaptation and levels of reference. *Journal of Vision*, 4, 693–710.
- Thompson, B. (2006). Colour constancy and Russellian representationalism. *Australasian Journal of Philosophy*, 84, 75–94.
- Thompson, B. (2007). Shoemaker on phenomenal content. *Philosophical Studies*, 135, 307–334.
- Thompson, B. (2008). Representationalism and the conceivability of inverted spectra. *Synthese*, 160, 203–213.
- Thompson, B. (2009). Senses for senses. *Australasian Journal of Philosophy*, 87, 99–117.
- Uchikawa, K., Uchikawa, H., & Boynton, R. (1989). Partial color constancy of isolated surface colors examined by a color-naming method. *Perception*, 18, 83–91.