

Perceptual Variation and Structuralism

JOHN MORRISON

Barnard College, Columbia University

Abstract

I use an old challenge to motivate a new view. The old challenge is due to variation in our perceptions of secondary qualities. The challenge is to say whose perceptions are accurate. The new view is about how we manage to perceive secondary qualities, and thus manage to perceive them accurately or inaccurately. I call it **perceptual structuralism**. I first introduce the challenge and point out drawbacks with traditional responses. I spend the rest of the paper motivating and defending a structuralist response. While I focus on color, both the challenge and the view generalize to the other secondary qualities.

1. Perceptual Variation

Our perceptual experiences tell us that lemons are yellow and sour. It's natural to infer that lemons *really are* yellow and sour. But perceptions vary, even among ordinary observers. Some perceive lemons as slightly greener and sweeter. This gives rise to a challenge that forces us to rethink the nature of perception and the objectivity of what we perceive.

For concreteness, I'll focus on the perceptions of two ordinary observers, Miriam and Aaron. I'll also focus on their perceptions of a particular lemon's color. These restrictions will make the challenge easier to understand and the responses easier to compare. But the challenge and responses are perfectly general. Toward the end of the paper I'll return to sourness and the other secondary qualities.

When Miriam and Aaron look at the same lemon, their perceptions differ. I think the best way to convey how their perceptions differ is to use your perceptions as a reference point. Let *phenomenal-greenish-yellow* be the phenomenal character of your perception when you report that a surface looks greenish yellow, and *phenomenal-pure-yellow* be the phenomenal character of your perception when you report that a surface looks unmixed, pure yellow under the same viewing conditions. Suppose that Aaron's perception of the lemon is phenomenal-greenish-yellow, and Miriam's perception of the lemon is phenomenal-pure-yellow. Who is accurately perceiving the lemon's shade, Miriam or Aaron?

There are four responses: neither, one, both, and indeterminate. These responses are traditionally developed: **neither** of them is accurately perceiving the lemon's shade, because the lemon isn't colored; **one** of them is accurately perceiving the lemon's shade, because they're perceiving incompatible shades, such as pure

yellow and greenish yellow; **both** of them are accurately perceiving the lemon's shade, because they're perceiving compatible shades, such as pure-yellow-for-Miriam and greenish-yellow-for-Aaron; and it's **indeterminate** who is accurately perceiving the lemon's shade, because it's indeterminate which shades Miriam and Aaron are perceiving. When developed in these ways, all four responses have significant drawbacks. I'll start by describing those drawbacks. I'll then use a theory of perception that I call **perceptual structuralism** to explain why it's possible for both Miriam and Aaron to accurately perceive the lemon's shade. In this way, I hope not only to suggest a new way of responding to the challenge of perceptual variation, but also to motivate a new view of perception. As we'll see, it is a significant departure from the views of Dretske (1981), Fodor (1987), and Millikan (1984), in part because it gives phenomenal characters a different role.

The challenge of perceptual variation isn't the only motivation for perceptual structuralism. In an earlier paper I argued that it gives us an attractive solution to a puzzle involving a series of pairwise indistinguishable chips (Morrison 2013). In an even earlier paper I identified a puzzle for anyone who accepts color realism and a traditional view of perception (Morrison 2012). At the time I thought that puzzle motivated color anti-realism, but I now think structuralism is a better solution.

Let's start by clarifying the challenge of perceptual variation. The challenge is to say whose perception is *accurate*. That's potentially confusing, because 'accurate' is equivocal. In one sense, to say that someone is more accurately perceiving an object is to say that their perception is more informative. If we were using 'accurate' in this way, it might be tempting to conclude that neither Miriam nor Aaron is accurately perceiving the lemon, because pigeons and other animals perceive more specific shades, and therefore have more informative perceptions. But that's not how I'm using 'accurate'. As I'm using it, Miriam and Aaron might still be accurately perceiving the lemon even if other animals perceive more specific shades. This is the sense of 'accurate' in which 'Brooklyn is in New York' is a perfectly accurate description of Brooklyn's location, because Brooklyn really is in New York, even though 'Brooklyn is in eastern New York' would have been more informative.

Next, the challenge is to choose between the perceptions of two *ordinary* perceivers. An ordinary perceiver is someone who passes all the standard tests of color acuity, such as the Farnsworth-Munsell 100 hue test. How do we know that there's variation among these perceivers? Let's review some of the compelling behavioral and physiological evidence. Suppose we let subjects turn a knob that controls the color of an image on a screen. The image originally looks greenish yellow (i.e., chartreuse), but, if they turn the knob far enough, it will look reddish yellow (i.e., orange). We ask subjects to turn the knob until the image looks pure, unmixed yellow, without any hint of green or red. As before, the same subject will reliably turn the knob to the same position. But different subjects will reliably turn the knob to different positions. Moreover, if we show subjects the images chosen by other subjects, they'll often report that the images look slightly greenish or reddish. Thus, even among ordinary perceivers, there's disagreement about which images are pure, unmixed yellow (for a summary of the data, see Kuehni 2004).

The variation in subjects' performance isn't random. Performance correlates with age and sex, though there's variation even within these groups (Scheffrin and Werner 1990; Neitz and Jacobs 1986). This makes it unlikely that the variation is due to different understandings of the terms 'pure', 'mixed', and 'same'. It is far more likely due to differences in the phenomenal characters of their perceptions. The fact that performance correlates with age and sex also makes it even harder to say whose perception is most accurate, because we can't uncritically assume that people of a certain age and gender have the most accurate perceptions (Block 1999, p.43).

This is compelling *behavioral* evidence that perceptual variation is pervasive. There's also compelling *physiological* evidence, including: the detectors in the eye most responsible for color perception vary in their sensitivities to light at different wavelengths (see Wyszecki and Stiles 1982, Ch 5; Stockman and Sharp 1999; Smith and Pokorny 2003, p.120–124). For example, the L-cones in one observer might on average be most sensitive to light at 570 nm, while the L-cones in another observer might on average be most sensitive to light at 565 nm. Also, the L-cones in one observer might be more sensitive in general than the L-cones in another observer, because the first observer's cones contain more of the photopigments responsible for generating the electric signals sent to the brain. These physiological differences help explain why color perception varies with age and sex. People in the same category tend to have more similar eyes.

The challenge of perceptual variation is one of philosophy's oldest. But Democritus, Heraclitus, Sextus, Lucretius, Plutarch, etc., formulated it differently. They asked us to choose between the perceptions of a healthy person and someone suffering from a disease.¹ This led to debates about whether the mere fact that a person is healthier is evidence that they perceive colors more accurately.² We don't need to engage with that debate, because we're trying to choose between the perceptions of two *healthy* observers. More recently, some have asked us to choose between the perceptions of an actual person and someone in a possible world where our eyes and brains are configured differently. This has led to debates about whether the mere fact that a person is actual is evidence that they perceive colors more accurately.³ We don't need to engage with that debate, because we're trying to choose between the perceptions of two *actual* observers.

Here's the plan: I'll first consider the traditional responses to this challenge (Sections 2–5). I'll then introduce structuralism and develop a structuralist response (Section 6).

2. Neither

The first response is that neither Aaron's nor Miriam's perception is accurate, because the lemon isn't colored. Given that the same challenge arises for almost any object, this would push one to deny that anything is colored. But we naturally think and talk about the world as colored; we naturally describe lemons as yellow, and we believe that a shirt reflects more sunlight because it's white rather than black. This isn't a decisive consideration. We should be prepared to revise our

natural ways of thinking and talking in response to compelling philosophical or empirical arguments. But, all other things being equal, I think we should preserve them.

Perceptual variation also seems like the wrong kind of reason to deny that anything is colored. Consider thermometers. Even though one thermometer measures a room's temperature as eighty degrees Fahrenheit and another thermometer measures the room's temperature as eighty-one degrees Fahrenheit, that isn't a reason to deny the room has a temperature. It's a reason to think that at least one of the thermometers is miscalibrated. Likewise, perceptual variation *by itself* doesn't seem like a reason to deny that anything is really colored. We'd need to be convinced that there are no plausible alternatives, and we're trying to develop such an alternative.

3. One

The second response is that at most one of their perceptions is accurate. I will call this **one-ism**. There are many ways to be a one-ist. A reflectance physicalist might think that at most one of Miriam's and Aaron's perceptions can be accurate because they're perceiving different reflectances, and the lemon reflects light in at most one of these ways. A dispositionalist might think that at most one of their perceptions can be accurate because they're perceiving dispositions to cause different perceptions in a certain kind of observer in a certain kind of context, and the lemon causes at most one of those perceptions in the relevant kind of observer in the relevant kind of context. A realist primitivist might think that they're perceiving different non-physical properties, and the lemon can instantiate at most one of them.

Regardless of how one-ism is developed, there's a problem. The problem is that, even if Miriam's perception is accurate, it's unclear what evidence could justify our belief that Miriam's perception is accurate. Unlike other cases of ignorance (of say, perhaps, events in the distant past), it's also unclear what could satisfactorily explain our ignorance. I think this is an unacceptable kind of ignorance, and should lead us to reject one-ism. In another paper, I develop this problem in full detail (Morrison manuscript b). In this paper, I'll just try to convey the general idea, because I suspect that most readers won't need all the details to be convinced of the problem, and would prefer to spend more time developing structuralism.

I'll begin by surveying two kinds of evidence that might seem to indicate that Miriam's perception is accurate. Keep in mind that we're not asking: What could *make it the case* that Miriam's perception is more accurate? That's a metaphysical question. We're asking an epistemological question, namely: What could *justify our belief* that Miriam's perception is more accurate?

The first kind of evidence is about the physiological differences between Aaron's and Miriam's eyes. Miriam's L-cones might on average be most sensitive to light at 570 nm, while Aaron's L-cones might on average be most sensitive to light at 565 nm. Or, Miriam's eyes might have proportionally more L-cones than M-cones. Or, Miriam's cones might be more sensitive to light than Aaron's detectors. And so on. However, it's unclear how this evidence alone could justify

our belief that Miriam's perception is more accurate. We'd still need to determine that people with eyes like Miriam's perceive colors more accurately. In other domains, this is easier to determine. If two microscopes produce different images, and one of them has a scratched lens, then we know which microscope is defective. But it's not similarly obvious why Aaron's eyes, rather than Miriam's eyes, are defective.

The second kind of evidence is about how the lemon reflects light. This evidence might seem especially helpful if you think that colors are reflectances. However, even in that case, it's unclear how this evidence could help justify our belief that Miriam's perception is more accurate. Even if we know the reflectance of the lemon, we're left wondering: Who is perceiving *that reflectance*? Is it Aaron, in virtue of his phenomenal-greenish-yellow perception, or is it Miriam, in virtue of her phenomenal-pure-yellow perception? If colors aren't reflectances, our situation is unchanged. For example, if colors are primitive properties, we're left wondering: Who is perceiving the color that *supervenes* on that reflectance?

There are other kinds of evidence that might seem helpful, including evidence about the phenomenal characters of Miriam's and Aaron's perceptions, the perceptions of others in their community, their shared evolutionary history, and even first-personal evidence about one's own perception. I consider all of this evidence in the other paper (Morrison manuscript b). But I hope that I've already said enough to convince most readers that our current and future evidence is unlikely to justify our belief that Miriam's perception is accurate.

Suppose I'm right, and that we shouldn't expect to justifiably that Miriam's perception is more accurate. So what? There's a lot we don't know. We don't know the current number of stars in distant clusters, because light from those clusters won't reach our telescopes for billions of years. We don't know the aggregate weight of all the chocolate in existence, because it's created and consumed too quickly. We don't know Socrates's exact height when he drank hemlock, because his corpse decomposed long ago. Given that we're ignorant of so much, what's wrong with concluding that we're ignorant about whose perception is accurate? This, in essence, is how color one-ists such as Stroud, Tye, Byrne, Hilbert, and Allen respond when asked to identify the person most accurately perceiving an object's color.⁴

But some kinds of ignorance are more acceptable than others. I just described cases in which there are identifiable causal processes, such as decomposition, preventing us from collecting the relevant information. Our ignorance about Miriam's perception isn't like that. Even if we had perfect access to all of the relevant evidence, we still wouldn't know whose perception is accurate.

This puts considerable pressure on one-ism. Consider motion. At some point in the earth's orbit, it moves 50 km/s relative to Mars, and 30 km/s relative to the sun. It's natural to assume that, in addition to these relative velocities, the earth also has an absolute velocity, that is, a velocity that isn't relative to any other bodies or frames of reference. However, if it does, we can't know it, due to the Special Theory of Relativity (see Bell 2004, p.75–76). To many philosophers and physicists, this seems like a sufficient reason to reject the assumption responsible for our

ignorance, namely that objects have absolute velocities (see Ismael and van Fraassen 2003, Maudlin 2012, p.67–68, Dasgupta 2016). I think that our ignorance about Miriam's perception would be similarly unacceptable, and therefore we should reject the assumption responsible for our ignorance, namely one-ism.

There's of course more to say about this objection. But hopefully I've said enough to motivate the search for another response.

4. Indeterminate

The third response is that it's indeterminate whose perception is accurate, because it's indeterminate what color Miriam is perceiving and it's indeterminate what color Aaron is perceiving. This response is hard to evaluate, because philosophers disagree about the nature of indeterminacy. But there's a general problem: this answer succeeds only if there's *enough* indeterminacy so that it's *always* indeterminate whose perception is accurate. If there's insufficient indeterminacy, the challenge is merely relocated to another case. Why think there's *enough* indeterminacy?

An example might help clarify the force of this question. Suppose we're looking at a stick and I claim "That stick is between nine and ten centimeters" and you claim "That stick is between ten and eleven centimeters." If the stick is a certain length, then, due to indeterminacy in the meaning of 'centimeter', it might be indeterminate whose claim is accurate. But if the stick were a little shorter, it would be determinate that my claim is accurate, and if the stick were a little longer, it would be determinate that your claim is accurate. In other words, even if the meaning of 'centimeter' is indeterminate, there isn't enough indeterminacy that it's *always* indeterminate whose claim is accurate. Why think there's enough indeterminacy in the case of color perception?

One might take inspiration from the philosophy of language. Because people use 'centimeter', 'bald', and 'heap' in different ways, some philosophers infer that it's indeterminate what these terms represent. It might be tempting to give a parallel explanation of why it's indeterminate what Miriam and Aaron perceive. In particular, it might be tempting to say that whenever there is variation in how two normal people perceive an object, it is indeterminate who is perceiving that object accurately. In that case, the amount of indeterminacy would be proportional to the amount of variation. This would be a tidy solution. But while what 'centimeter', 'bald', and 'heap' represent might depend on how others use these terms, and thus while variation in how others use these terms might give rise to indeterminacy, what we perceive doesn't seem to depend on contingent facts about what other people perceive. If everyone with eyes and brains unlike Miriam's died tomorrow, that wouldn't change what property she perceives. It also wouldn't make a difference if they died shortly before her birth, or even never existed. I think this reflects a fundamental difference in the kinds of intentionality involved in perception and linguistic communication. When we speak, we intend to use terms as others use them, but when we perceive, our perception doesn't include an intention to see what others see. Perception involves a less sophisticated kind of intentionality, at least in this respect. Thus, while what our terms represent might depend on how others use

them, what we perceive doesn't seem to depend on what others perceive. For this reason, I don't think that perceptual variation can be responsible for indeterminacy in what we perceive, and I'm not sure what else could explain why there's enough indeterminacy.

This brings out an important difference between the challenge of saying whose *perception* is accurate and the challenge of saying whose *sentence* is accurate, after Miriam asserts "The lemon is pure yellow," and Aaron counters "The lemon is greenish yellow." The meanings of these sentences arguably depend on how other people use 'pure yellow' and 'greenish yellow'. Thus, if there's widespread inconsistency in how people in our community would describe the lemon, then it might be indeterminate whose assertion is accurate.

Of course, one needn't think about the relevant indeterminacy as indeterminacy in what Miriam and Aaron perceive. For example, it might instead be indeterminate what property the lemon *instantiates*. In particular, it might be indeterminate whether it instantiates pure yellow or greenish yellow. Nonetheless, the problem remains. We're still left wondering: Why think there's enough indeterminacy so that it's always indeterminate whose perception is accurate?

As noted above, this is a difficult response to evaluate, because of disagreements about the nature of indeterminacy. Someone might be able to develop a way of thinking about indeterminacy that avoids the problem. But, even if there is such a way of thinking, I think we should still prefer views that imply that people can accurately perceive the colors of objects even when there's interpersonal variation. I'll explain why in the next section.

5. Both

The fourth and final response is that both of their perceptions are accurate. I will call this **both-ism**. I think there's something right about this response. Suppose we fill one hundred unmarked tubes with different amounts of mercury and then place them in a warm room. Their mercury will rise to different levels. It would be silly to then ask which tube of all the tubes is *accurately* measuring the room's temperature. Each tube just registers the temperature in a different way. Likewise, nature has filled our eyes with different amounts of the relevant kinds of detectors, distributed those detectors in different patterns along our retinas, wired our brains to respond differently to the signals sent from our eyes, and so on. Consequently, people's eyes and brains respond differently to the same object, producing different perceptions. As with the unmarked tubes of mercury, it seems misguided to ask which perceptual system of all perceptual systems is accurately measuring the lemon's color. Different perceptual systems just register the object's surface in different ways. For this reason, I think there's something right about views according to which Miriam's perception and Aaron's perception can both be accurate.

But I think there's something wrong about how both-ism is traditionally developed. In particular, I don't think it should be combined with a widespread view about the nature of perception that I call "perceptual atomism." Let's introduce perceptual atomism by elaborating on our initial example. If Miriam is perceiving

the lemon alongside a lime, she doesn't just perceive the lemon as pure yellow. She also perceives it as yellower than the lime. Let's separate these elements:

1. Miriam perceives the lemon as pure yellow.
2. Miriam perceives the lemon as yellower than the lime.

What explains (1) and (2)? We're looking for a kind of explanation that's familiar from the philosophy of language. Philosophers of language want to explain how 'Napoleon' manages to refer to Napoleon. Some appeal to causal relations between 'Napoleon' and Napoleon. Others appeal to descriptions associated with 'Napoleon' that single out Napoleon (e.g., the French general who lost at Waterloo in 1815). These aren't the only explanations of how 'Napoleon' manages to refer to Napoleon, but I hope they're enough to get a grip on the relevant kind of explanation. Explanations of this kind are called "metasemantic" because they don't just indicate the objects and properties that our terms refer to (that's semantics); they explain how our terms manage to refer to those objects and properties.

We're searching for a similar kind of explanation. We want to explain how Miriam manages to perceive the lemon as pure yellow, and the lemon as yellower than the lime. Explanations of this kind are often called "psychosemantic." They don't just indicate the objects and properties that our psychological states refer to; they explain how our psychological states manage to refer to those objects and properties.

Perceptual atomists start by explaining facts like (1) by appealing to relations involving perceptions of the same kind. They typically group perceptions together according to their phenomenal characters, so that two perceptions belong to the same kind if they have the same phenomenal character. What relations do they appeal to? They often appeal to *causal* relations. For example, an atomist might say that Miriam perceives the lemon as pure yellow because pure yellow objects cause that kind of perception in normal perceivers under normal conditions (Peacocke 1984, p.373), or in ideal perceivers under ideal conditions (Tye 2002, p.138). Alternatively, an atomist might say that Miriam perceives the lemon as pure yellow because of a *primitive, non-causal* relation between that kind of perception and pure yellowness (Pautz 2010a, p.58–60). Regardless, perceptions of other kinds, including perceptions of green limes and orange pumpkins, don't play a role. Each kind of perception is explanatorily independent of the others.


Perceptual atomists then explain facts like (2) by appealing to facts like (1). Atomists claim that Miriam perceives the lemon as yellower than the lime because she perceives the lemon as pure yellow and the lime as pure green. They think of color perception like a by-the-numbers painting, where what's explanatorily basic are assignments of colors to individual objects, such as pure yellow to the lemon and pure green to the lime. This isn't the whole story, but the details needn't concern us (for the details, see my manuscript a.) For our purposes, it's enough that facts like (1) are supposed to be part of the explanation for facts like (2).

Why call this "perceptual atomism"? I think there's a helpful parallel with metaphysical atomism. Metaphysical atomists claim that the lemon and lime result from

combining more fundamental objects that are metaphysically independent of one another. At a minimum, the most fundamental objects are independent of each other in that we can say what each is without mentioning the others. They might also be independent of each other in that it's possible for each to exist without the others. Likewise, perceptual atomists claim that Miriam's total perception of the lemon and lime results from combining more basic perceptions, namely her perception of the lemon as yellow and her perception of the lime as green, and that these perceptions are explanatorily independent of one another. At a minimum, the perceptual atomists' explanations are independent of each other in that her explanation of each doesn't mention the others. They might also be independent in that it's possible for her to perceive the lemon as yellow even if she's unable to perceive anything as green.

Both-ism is traditionally combined with atomism. But combining both-ism and atomism makes it difficult to explain why our perceptions are sometimes inaccurate. I develop this problem in full detail elsewhere (manuscript a). In this paper, I'll once again just try to convey the general idea.

We've been focusing on variation in what different people perceive. But there's also variation in what the same person perceives. Consider Kitaoka's (2006) lightness illusion:



Perhaps surprisingly, the left and right squares are intrinsically alike. More precisely, they reflect light in the same way, and, as a result, could be interchanged without affecting the illusion. Which square, if either, are you accurately perceiving?

An atomist shouldn't respond that you're accurately perceiving only *one* of the squares. To see why, suppose that you're accurately perceiving only the left square. The problem is that, from an atomistic point of view, nothing seems to justify the belief that you're accurately perceiving only the left square. Your perception of the left square is phenomenal-dark-gray, while your perception of the right square is phenomenal-medium-gray. But there's no reason to believe that the accurate perception is phenomenal-dark-gray rather than phenomenal-medium-gray. Similarly, your perception of the left square is the result of viewing it against a relatively light background, while your perception of the right square is the result of viewing it against a relatively dark background. But there's no reason to believe that the accurate perception results from viewing the left square against a lighter

background rather than a darker background. More generally, none of the intrinsic features of these perceptions, including their phenomenal characters, and none of their relations to the squares, including their causal relations, seem to justify the belief that you're accurately perceiving only the left square. From an atomistic point of view, it's unclear what else could justify that belief, because atomists say that only your perception's intrinsic features and relations to objects are relevant. Thus, an atomist who claims that you're accurately perceiving only one of the squares is apparently committed to ignorance about which square it is. And, like before, there isn't the right kind of causal explanation of why we can't identify the relevant square.

For this reason, I think that atomists should respond that you're accurately perceiving *both* squares.⁵ But how is that possible? That is: How does your perception of the left square manage to represent its color, and how does your perception of the right square manage to represent its color? I think that atomists should say that it's due to causal relations in your *current viewing contexts*. More precisely, I think that atomists should say that you represent the left square's color because objects with that color cause phenomenal-dark-gray perceptions in you in the left context, and you represent the right square's color because objects with that color cause phenomenal-medium-gray perceptions in you in the right context. It would take a long time to properly defend this recommendation (see my manuscript a). But its appeal is straightforward: it would explain why your perceptions of both squares are accurate; it wouldn't lead to ignorance about the relevant context or contexts; it wouldn't require us to attribute properties to the squares that go beyond what's attributed by our best scientific theories; and it wouldn't lead to ignorance about whether you stand in the relevant relations to the squares' colors, as it might if primitive relations were instead relevant.

There are two proposals along these lines. The first is about the colors themselves. The second is about the descriptions that pick out the colors. More specifically:

The first proposal is that you're perceiving the left square as instantiating *causes phenomenal-dark-gray perceptions in me in the left context* and you're perceiving the right square as *causes phenomenal-medium-gray perceptions in me in the right context*. The left square has the first property, and the right square has the second property. According to this proposal, these properties are colors. In that case, both of your color perceptions are accurate. Cohen (2004, 2009) endorses and develops this proposal. It's the proposal about the color themselves.

The second proposal is that you're perceiving the left square as instantiating whatever property satisfies the description: the property that disposes the left square to cause phenomenal-dark-gray perceptions in me in the left context. Likewise, you're perceiving the right square as instantiating whatever property satisfies the description: the property that disposes the right square to cause phenomenal-dark-gray perceptions in me in the right context. Because the squares are intrinsically alike, it's possible that the same property satisfies both descriptions. But, regardless, these descriptions pick out properties of the left and right squares. Thus, you're accurately perceiving both squares. Moreover, you're accurately perceiving *the colors* of both squares, because to be a color is just to satisfy a description of this kind.

Jackson and Pargetter (1987) and McLaughlin (2003) endorse and develop this proposal.

The problem with both proposals is that almost all color perceptions would be accurate. Consider the first proposal. If you have a phenomenal-dark-gray perception in a context, and your perception is caused by an external object, that object instantiates *causes phenomenally-dark-gray perceptions in you in that context*. You also perceive an object as instantiating *causes phenomenally-dark-gray perceptions in me in that context*, because that's what you perceive when you have a phenomenally-dark-gray perception in that context. Thus, unless you're misidentifying the object causing your perception, your perception is accurate. More generally, almost all your perceptions are accurate, including perceptions we'd like to characterize as illusory. The second proposal has the same consequence, with the unimportant difference that your perception represents the property that *disposes* objects to cause a phenomenally-dark-gray perception. Yet we sometimes seem to misperceive the colors of objects, even under normal conditions. In Kitaoka's illusion, we perceive the left square as darker, even though the squares are the same, and thus seem to be misperceiving at least one of the squares. Of course, there's more to say about this problem (see my manuscript a). But I hope this is enough for present purposes.

Cohen acknowledges this difficulty (see also Clark 2000, p.226–227). He tries to mitigate it by claiming that standard examples of color illusion involve false predictions about how an object will look in other contexts. With respect to Kitaoka's illusion, you might falsely predict that the squares would look different if viewed by certain perceivers in certain contexts. Which perceivers? Which contexts? According to Cohen, that's determined by our linguistic community, because illusions involve the misapplication of color predicates, and our the linguistic community uses certain perceivers and contexts to establish the correct application of predicates such as 'gray' and 'yellow'. He thereby relocates color illusions to thought and language. This is unsatisfying, because it implies that animals are incapable of color illusions, unless they cognitively single out certain perceivers and contexts as normal (Cohen 2007, p.338–339 acknowledges this consequence). It also implies that whether a perception is illusory can change as a result of a change in our language, or perhaps even just a change in which language we're inclined to use at that moment, as long as users of those languages single out different perceivers and contexts as normal. Perhaps most fundamentally, this proposal is unsatisfying because our errors are sometimes entirely perceptual. In Kitaoka's illusion, we're *misperceiving* the left square as darker, and thus we're *misperceiving* at least one of the squares. Our error is perceptual, not cognitive or linguistic. There is disagreement about where to draw the border between perception and cognition. But I think anyone who draws the border so that these errors are non-perceptual is drawing the border in the wrong place.

Let's summarize: From an atomistic point of view, there's no way to justify the belief that you're accurately perceiving only the left square. This pushes atomists to say that you're accurately perceiving both squares, because saying that you're accurately perceiving only one square would lead to unacceptable ignorance. But

the best proposals imply that almost all of your perceptions are accurate. As we'll see in the next section, structuralists aren't pushed to the same conclusion, because they can explain why you're misperceiving one square but not the other, without committing us to an unacceptable kind of ignorance.

6. Structuralism

The alternative I prefer is **perceptual structuralism**. It isn't the only alternative to atomism.⁶ But I think it's the best. Structuralists explain facts like (2) by appealing to *all* of Miriam's past perceptions, including her perceptions of green limes and orange pumpkins. Then, on that basis, they explain facts like (1). Structuralists thereby reverse the traditional, atomistic order of explanation.

As I think the view should be developed, there's a helpful parallel between perceptual systems and thermometers. Suppose Ruth doesn't know anything about thermometers. In fact, suppose she's incapable of direct experiences of temperature—nothing feels warm or cool to her. One day we give her an unmarked tube partially filled with mercury and ask her to wander around a drafty old mansion with rooms of varying temperatures. When she first looks at her thermometer, the mercury level won't be a meaningful source of information. For all she knows, it might be a fixed, unmovable feature of the tube. But after she walks into a different room and watches the mercury level rise, she can compare the rooms, because she can rank them by the level of mercury in her tube. Suppose she uses 'hotness', a word she invented, to rank them. She might say, "This room is hotter than the last room." As she visits more rooms, she can make more and more complicated comparisons, because her ranking will include many more rooms. Over time, she might even develop a sense of how a given measurement relates to measurements she's made in the past, allowing her to make comparisons to past rooms long after she's forgotten most of their details. She might say of a given room, "This room is slightly hotter than any of the rooms I've ever visited." Her comparisons will eventually yield a ranking that matches differences in temperature better than any other differences. When that happens, her claim that one room is hotter than another room will be a comparison of their temperatures. Moreover, her comparisons will describe each new room as having a temperature in an approximate range, such as 20–25°C-ish. That is, by ranking a room relative to other rooms, she'll thereby describe it as having a certain thermal property. Nonetheless, if she doesn't have a scientific background, she won't be able to tell us anything substantive about that property, or, more generally, what it is for a room to be hotter than another room.

I think that perceptual systems play a role similar to Ruth's thermometer. This isn't to deny there are important differences, including that Ruth's inferences are far more deliberate than what's involved in perception. But Ruth's thermometer is still a helpful model. Here's the rough idea: When Miriam perceives objects, she uses the phenomenal characters of her experiences to perceive some of those objects as yellower, greener, bluer, redder, whiter, or blacker than others. Her perceptions eventually yield a ranking that matches certain differences

between the surfaces of those objects better than others. For concreteness, let's suppose they match differences in how those objects reflect, emit, and transmit light (hereafter just: reflect light). Her perception of a lemon as yellower than a lime is then a perception of a certain kind of difference in how they reflect light. Of course, her perception of a lemon isn't normally exhausted by her perception of it as yellower than a single lime. She also perceives it as yellower than other objects. Combined, these perceptions "describe" the lemon as reflecting light in a certain kind of way. Since we're assuming that colors are ways of reflecting light, her perceptions are thereby describing the lemon's color. Nonetheless, without a scientific background she couldn't say anything substantive about that property, or, more generally, what it is for an object to be yellower than another object.

On this way of thinking, phenomenal characters are like meaningless symbols that become meaningful over time, as the result of comparisons. Perception is a *gradual* and *individual* achievement.

Structuralists can explain why it's possible for both Miriam and Aaron to accurately perceive the lemon's shade. The details will emerge gradually in the subsections that follow. But, once again, here's the rough idea: Structuralists can start by noting that Aaron perceives the lemon as slightly greener than other objects, but Miriam doesn't. As a result, Aaron's perceptions will eventually match one kind of difference in how objects reflect light, making that the kind of difference he perceives between the lemon and the lime, while Miriam's perceptions will eventually match another kind of difference, making that the kind of difference she perceives between the lemon and the lime. Compared to Aaron, Miriam might perceive a kind of difference that places more weight on the proportion of light reflected at certain wavelengths, for example. Nonetheless, both might be accurately perceiving a kind of difference that really exists between the lemon and lime. Moreover, the relations they perceive between the lemon, lime, and other objects might together describe how the lemon reflects light. In that case, their perceptions of the lemon's color are both accurate. Structuralists can therefore explain how it is possible for Miriam and Aaron to both accurately perceive the color of the lemon. An imperfect model involving absolute locations might help: even if Miriam says the lemon is one inch to the left of the lime, and Aaron says the lemon is two inches to the right of a tangerine, and thus they're describing different relations involving different objects, it's possible for both of them to be accurately describing the lemon's absolute location.

Structuralism is a view about how we manage to successfully interact with our environment, and thus about how we manage to perceive the properties that objects actually instantiate. It is thus a view about what I will call the "environmental content" of our perceptions, also known as its "objective representation" (Kriegel 2013) and "ordinary content" (Chalmers 2006, p.70). Structuralism is noncommittal about whether our perceptions have any other kind of content. For example, it is noncommittal about whether our perceptions also represent what Chalmers (2006) calls "Edenic colors," which are uninstantiated properties that perfectly match our phenomenal characters. As I hope will become clear, I don't think that

structuralists need to appeal to Edenic colors to explain the environmental contents of our perceptions. But they are nonetheless free to do so, perhaps to help satisfy other explanatory demands, such as those imposed by Johnston (2004, p.130–131), Chalmers (2006, p.61–66), and Pautz (2010b, p.266–271).

Structuralism also isn't a view about the metaphysical nature of the colors. At least in principle, structuralists can identify yellow with a disposition to reflect light ("reflectance physicalism"), an intrinsic property of the lemon that isn't equivalent to any of its physical properties ("primitivism"), or a disposition to cause certain kinds of perceptions in certain kinds of observers ("dispositionalism"). Structuralism might be used to argue for one of these identifications, but it is not itself a view about the metaphysical nature of color. To the extent that we can describe structuralism as a metaphysical view at all, we should describe it as a view about the nature of perceptual representation, not the colors.

I will divide this section into eight subsections, each describing a different part of the view. I will describe the role and nature of phenomenal characters (6.1), the relations we perceive (6.2), the objects we perceive relations between (6.3), what it is for our perceptions to match relations in the external world (6.4), the relations involving light I think our perceptions match in the actual world (6.5), the colors we thereby perceive (6.6), why both Miriam and Aaron might be accurately perceiving the lemon's shade (6.7), and some of structuralism's notable features, including how it generalizes beyond color (6.8).

6.1 Phenomenal Characters

Let's use the story involving Ruth to structure our discussion. After Ruth walks into the second room, she might say, "The level of the mercury in my tube is higher." Because she's treating the mercury in her tube as a measurement, she might use it to compare the two rooms, leading her to report, "The magnitude I'm measuring is greater in this room." Or, as she might prefer to say, "This room is hotter." Thus, Ruth represents the mercury in her tube *as a measurement in her tube*, and she represents a room as hotter as the result of representing a change in mercury level.

According to structuralists, phenomenal characters are unlike Ruth's mercury levels in at least two respects, each corresponding to one of the less discussed mysteries of consciousness. The most discussed mystery is how brains give rise to consciousness at all. This mystery has rightfully been a central preoccupation of philosophers interested in consciousness. But two mysteries that are less discussed, or at least less general, are more relevant for our purposes.

The first is why we're conscious of external objects. Like a submarine captain who sees only the positions of his ship's gauges, and must infer the condition of the surrounding ocean, we might have been conscious only of the states of our own brain. Phenomenologists could have completely described consciousness with reports like "Neurons firing fast!" or "Lots of activity in the left half of the visual cortex!" Fortunately for most of us, though perhaps unfortunately for neuroscientists, our consciousness isn't like that. We're conscious of external objects, rather than our own brains. This general feature of consciousness extends to the brain's measurements of external objects. The brain measures the lemon's surface by combining inputs from

several kinds of detectors. The phenomenal characters of our perceptions are these measurements. But we're not aware of these phenomenal characters *as measurements in the brain*. In this respect, I think that the brain is like a thermometer that displays its measurement of the surrounding air by projecting a dimmer or brighter light onto that very atmosphere, or by broadcasting a softer or louder sound, except, of course, in the case of perception, nothing is really projected or broadcast. It's mysterious how brains can give rise to this kind of perceptual consciousness.

This feature of consciousness has an important consequence. Ruth is unlikely to confuse features of the mercury in her tube for features of the surrounding air. If the volume of the mercury in her tube is two cubic centimeters, she's unlikely to infer that the volume of the surrounding air is two cubic centimeters. In most cases, we're similarly unlikely to confuse features of our representations for features of what they represent. Even though 'cats' has four spatially distinct parts ('c', 'a', 't', 's'), we're unlikely to infer that cats have four spatially distinct parts. But with respect to perceptual experiences, this is a natural kind of confusion, because our awareness of phenomenal characters is directed outward, onto the very objects they measure. This confusion might be responsible for the strong initial pull of atomism. Our relation to individual phenomenal characters seems basic; we're not aware of phenomenal-pure-yellow because we're aware of other phenomenal characters. This might mislead us into thinking that our *perceptions* of individual colors must be basic. That is, the strong initial pull of atomism might result from confusing a feature of our relation to *a measurement in the brain* for a feature of our relation to *the external quantities it measures*.

The second mystery is how we represent external objects and their relations without first representing our own phenomenal characters. Ruth represents a difference in the relevant magnitude ("the first room was hotter") by first representing a difference in her thermometer's mercury level ("the level of the mercury when I was in the first room was higher"). Consciousness isn't like that. Miriam doesn't represent the lemon as yellower than the lime because she represents a difference in her phenomenal characters. Miriam's phenomenal characters are *responsible* for her comparisons between objects, but it's not the sense in which Ruth's mercury levels are responsible for her comparisons between rooms. This feature of consciousness is also manifest in perceptual demonstratives. Demonstratives like 'that property' and 'that difference' naturally pick out properties and relations of external objects, rather than properties and relations of phenomenal characters. Assuming it's even possible to use these demonstratives to pick out our own phenomenal characters, it takes a special kind of effort, and perhaps also a special kind of sophistication. It's mysterious what gives phenomenal characters this feature.

I hope we'll one day solve these mysteries. But we haven't yet. As a result, we can't yet give a fully satisfying explanation of *how* Miriam's phenomenal characters allow her to perceive relations between external objects. But I hope it's plausible enough that her phenomenal characters play that role.

Structuralism imposes at least one constraint on the nature of phenomenal characters: they must be properties of Miriam's mind, rather than properties of external objects. But beyond that, structuralism is fairly noncommittal. Some

examples: (i) It is noncommittal about how to individuate phenomenal characters. Structuralists can even agree with James (1950, Volume 1, p.231) that two perceptions never share the same phenomenal character. For the structuralist, all that's important is that Miriam's phenomenal characters allow her to perceive objects as yellower and greener, and mere similarities between her phenomenal characters should be sufficient. (ii) It is noncommittal about whether phenomenal characters are properties of Miriam's brain or an immaterial mind. (iii) It is noncommittal about whether individual phenomenal characters, such as phenomenal-pure-yellow, are metaphysically prior to total phenomenal characters, such as the total phenomenal character of your perception right now. (iv) It is noncommittal about the nature of the awareness relation. Structuralists might even reject what I said above, and agree with higher-order theorists that awareness is a kind of representation (see, e.g., Rosenthal 1986).

This is a partial list. But I hope it makes clear that structuralism can be combined with many different views about phenomenal characters.

6.2 Relations

Whether Ruth represents a room as hotter or colder is a *simple function* of her mercury level. If the mercury level rises, she represents a room as hotter, and if it drops, she represents a room as colder. Moreover, Ruth's mercury levels allow her to represent similarities and differences along only *one axis* of variation, because her mercury level only rises and falls. Finally, Ruth doesn't *perceive* rooms as hotter or colder, at least given some plausible assumptions about the distinction between perception and cognition. Instead, Ruth judges that a room is hotter or colder.

Miriam's representations of similarities and differences are unlike Ruth's representations in all three respects. First, Miriam doesn't represent an object as yellower, greener, etc., than another object just because there's a difference in her phenomenal characters. If Miriam looks at a pair of lemons, one directly illuminated by the sun and the other covered by a shadow, her phenomenal characters will differ, but she might still perceive the lemons as the same color. Likewise, if she watches a lemon under an intensifying light, her phenomenal characters will differ from one moment to the next, but she might still represent it as the same color. As these examples suggest, the process underlying her representations of relations between objects tries to discount for lighting, background, distance, viewing angle, and other contextual factors. The neural and computational details are complicated, and there's an enormous scientific literature devoted to sorting them out. For our purposes, those details won't matter. All that matters is that, while the default is to represent an object as yellower, greener, etc., whenever there's a difference in phenomenal character, that default is sometimes overridden after the underlying process takes contextual factors into account.⁷

Second, Miriam doesn't just represent relations between objects, she *perceives* them, at least given some plausible assumptions about the distinction between perception and cognition. To start, the underlying process is *automatic*. When Miriam looks at the lemon and lime, she represents the lemon as yellower, even if she didn't previously intend to compare them. Likewise, in Kitaoka's illusion, she perceives

the left square as darker than the right square, regardless of her previous intentions. The relevant process is also *fast*. From a first-personal point of view, Miriam can't detect any delay between her initial phenomenal characters and her representation of the lemon as yellower than the lime. Perhaps that's because these events are simultaneous. Or perhaps that's because there isn't a sufficient delay between them. A delay might be explained as follows: Processing in her early visual system might give rise to her initial phenomenal characters, without taking into account many contextual factors. Processing further downstream might represent two lemons as the same color, after taking into account contextual factors, such as differences in illumination. This downstream processing might even give rise to its own kind of phenomenal character, such as the kind of phenomenal character that accompanies the perception of patterns, including color patterns. Regardless, if there is a delay, it is too short for Miriam to notice. Finally, the relevant process is *dissociable* from what she knows, in that what she thereby represents can conflict with what she knows. Even if she knows that the squares in Kitaoka's illusion are interchangeable, she'll still represent the left square as darker. This doesn't imply that what she thereby represents is completely independent of what she knows. There might be interactions. It's just that, in many cases, such as Kitaoka's illusion, there's conflict. Because Miriam's representations are automatic, fast, dissociable from belief, and also directed at our nearby environment, I will describe them as *perceptions*.

Third, Miriam's phenomenal characters allow her to perceive similarities and differences along at least three axes of variation, because her phenomenal characters vary in at least three ways. Thanks to these three kinds of variation, objects can look redder or greener, yellower or bluer, and lighter or darker. These perceptions often involve more than mere rankings. They often involve degrees, as when she perceives a lemon as *significantly* yellower than a salmon steak. I'll later suggest that these degrees are grounded in ratios to maximally yellow objects, i.e., the objects we perceive as equally yellow if not yellower than all other objects. There's an interesting question of whether she perceives other kinds of relations, such as ratios between differences. For example, she might perceive that the difference between the squares on the left is more than twice the difference between the squares on the right:



Another interesting question is whether she perceives negative relations, such as that an object is *not yellower* than another. One might also ask which relations are perceptually primitive. When she perceives a grape as purpler than a lemon, this might be a perceptually primitive comparison, or it might consist in perceiving the grape as redder and bluer. Let's set these questions aside for another occasion.

6.3 Objects

Miriam perceives the lemon as yellower than other objects. Which other objects? We can divide them into three categories.

First, she perceives the lemon as yellower than objects she's currently perceiving. She might perceive the lemon as somewhat yellower than a grapefruit, significantly yellower than a lime, and just as yellow as another lemon. This leaves open whether she perceives relations between the lemon and *all* other objects she's currently perceiving, including those at the extremes of her visual field. I'm inclined to think not, but for the purposes of this paper, it won't matter.

Second, Miriam perceives the lemon as yellower than objects she perceived in the recent past. If Miriam is attending to an unripe lemon, and we replace it with a ripe lemon, she might perceive the new ripe lemon as yellower than the old unripe lemon. Likewise, if Miriam is viewing the lemon against an orange background that's fading to white, she might perceive the contrast between the lemon and its background as increasing, in that she perceives the difference between the lemon and its background as greater than before. She also perceives relations between the lemon as it is now and as it was a moment ago. If she's watching a time-lapse video of the lemon ripening, she'll perceive it as yellower and yellower. But if the lights dim under normal conditions, she might perceive the lemon as the same color as it was before; that is, she'll perceive its color as constant.

This might seem to have an unacceptable result. Consider the case in which the unripe lemon is replaced by the ripe lemon. What I said might seem to imply that Miriam is perceiving the old lemon *in exactly the same sense* that she's now perceiving the new lemon, and that's clearly absurd. But what I said doesn't have this consequence. Miriam might perceive the unripe lemon in a different sense. Husserl (1980) called this other kind of perception "retention." At least in this case, I think we should classify retention as a kind of perception, because it's fast, automatic, dissociable from our beliefs, and constitutive of our perception of present objects, and involves processing in the visual system. But if you have special reasons for using 'perception' more narrowly, our disagreement won't have any further repercussions, except that if you endorse structuralism you'll have to say that we perceive the lemon's color in virtue of what we both perceive and retain.

Third, in addition to perceiving the lemon as yellower than objects she's currently perceiving and recently perceived, she perceives it as yellower than objects she perceived in the more distant past. This might sound especially strange, so it's worth going into more detail. I'll start by listing a number of examples that suggest we often have an appreciation for how an object relates to objects we've perceived before. I'll then defend the claim that in such cases we're perceiving relations to those objects. I'll end by applying this general framework to color perception. As with many of the other claims that are constitutive of structuralism, I won't have decisive arguments for any of these claims. My more limited goal is to show that they're coherent and plausible.

Over time, I think we all develop an appreciation for how a given object relates to objects we've seen before. Object recognition is a good model. When Miriam recognizes Abraham, she appreciates that this is a person she's met before, though perhaps without consciously recalling any of the details. This appreciation might just consist in a feeling of familiarity, or at least the absence of a feeling of novelty.

If she doesn't know his name, she might just say something like, "him again." In addition to this minimal reaction, she might also have more sophisticated reactions. If he was cruel, she might become afraid. Or, if he was kind, she might become calm. She might anticipate the sound of his voice, so that she'd be surprised if she later heard a voice with a deeper timbre. She might subvocalize his name. She might even have a conscious, episodic memory of their first encounter. Underlying all of these reactions is a minimal appreciation that this is a person she's met before. Even animals and babies evince this kind of appreciation in their differential responses to parents, siblings, and strangers.

Object recognition isn't unique in this respect. You might recognize a fabric without consciously recalling any of the blouses that were made from it, a perfume without consciously recalling any of the people who were wearing it, and a spice without consciously recalling any of the dishes that were flavored by it. If you don't know the relevant labels, you might just say, "that again." Switching to kinds, you might recognize the breed of a dog, the style of a beer, or the formation of a cloud, without consciously recalling any other particular dogs, beers, or clouds. If you don't know the relevant labels, you might just say, "another of those."

In many cases, this minimal appreciation doesn't just indicate that you've encountered something like it before, but also includes an appreciation of how an object relates to things of related kinds. Your attention might be drawn to a blouse that's especially lustrous, a perfume that's especially floral, a spice that's especially pungent, a dog that's especially mangy, a beer that's especially hoppy, or a cloud that's especially wispy. This shift in attention is due to your appreciation of how the object you're currently perceiving compares to blouses, perfumes, spices, dogs, beers, or clouds you've perceived before. Once again, animals and babies evince this kind of appreciation in their learned reactions to certain kinds of stimuli, and calibrated reactions to similar but novel kinds of stimuli.

Suppose I'm right and we often appreciate how an object relates to objects we've perceived before. What, if anything, do we perceive in such cases? I claim that we perceive *a relation to those objects*. This might sound strange, because more traditional views don't have this implication.

Once again, object recognition is a good model. Suppose you look at a person and have the kind of reaction best expressed by "him again." A traditional view is that you perceive that person *as Abraham*, because he's the person who is *disposed* to put you in the kind of perceptual state that produces that reaction. According to this view, you just perceive Abraham, a particular object. The alternative I prefer is that you perceive him *as the same person* as whomever *actually* caused that kind of perceptual state before. If Abraham is the person that best fits that description, you thereby perceive Abraham. But what's most fundamental is your perception of sameness, a relation. There's a lot more to say about this account of object recognition, but let's instead turn to other kinds of recognition, because they provide better models for our account of color perception.

Suppose you see a dog and have an appreciation for how it relates to dogs you've seen before. A traditional view is that you perceive the dog as instantiating the *property* shared by all the objects that are *disposed* to cause the relevant kind of

perception, and that you perceive similarities and differences between that property and the *properties* shared by objects disposed to cause relevantly similar perceptions. You might perceive the dog as instantiating *is a poodle* and you might perceive a dissimilarity between the property *is a poodle* and the property *is a rottweiler*. According to this kind of view, you perceive the dog you're seeing, a property of that dog, and perhaps also a relation between that property and other properties. You might even perceive the locations of these properties in what's called a property space. Thus, if you're drawn to this kind of view, it will sound strange to hear people say that you perceive a relation between the dog in front of you and dogs you're seen before. Nonetheless, this is an implication of the view I prefer. According to this alternative view, you perceive the dog as belonging to the *group* of dogs that have *actually* caused the relevant kind of perception, and you perceive it as more or less similar to objects in *other groups*. You might thereby perceive the dog as a poodle, if that's the breed of dog that best fits the relevant descriptions. But what's fundamental is your perception of the relation between the dog you're currently seeing and the dogs you've seen before.

Several additional points might soften initial resistance to this view. To start, I'm not saying you perceive each individual dog in the relevant groups. There's a helpful parallel with how some nominalists about properties will analyze your claim that "Socrates is white." According to these philosophers, Socrates is white just in case he's a member of the class of all white objects (for discussion, see Armstrong 1978, p.28–34). Nobody should think that these philosophers are thereby committed to the claim that when you say that Socrates is white, you're *individually* representing all of the other objects in that class. Likewise, nobody should object to structuralism that it implies that we're *individually* perceiving each of the objects we've perceived before. These are instances of plural reference, as when you say, "the dogs are surrounding the pond," thereby describing a group without describing each individual in the group. You can represent a dog as *one of them* without representing all the other dogs in that group.

I'm also not saying that you *perceive* the relevant groups *in the same sense* in which you're perceiving the lemon in front of you, or even in the sense that you're perceptually retaining a lemon that was just removed. Object recognition might again be helpful. When you perceptually recognize an object, you're relating it to what you've perceived before, without consciously recalling that previous encounter. Similarly, the property nominalist mentioned earlier isn't committed to the view than when you say, "Socrates is white," you're referring to Socrates in the same sense that you're referring to the people you're disposed to call 'white'. In the case of perception, we might call your relation to those other objects 'grouping'. I think we should say you *perceive* the relevant groups, because this kind of representation is fast, automatic, dissociable from our beliefs, and constitutive of our perception of present objects, and involves processing in the visual system. But, as before, if you have special reasons for using 'perception' more narrowly, our disagreement won't have any further repercussions, except that if you accept structuralism you'll need to say that we perceive the lemon's color in virtue of perceiving, retaining, and grouping.

There's a lot more to say about this alternative view, and I'll fill in more details in the following pages. But this is enough to apply it to color perception. When Miriam perceives the lemon, she has an appreciation for how it relates to objects she's perceived before. In particular, she has an appreciation for how it compares to objects in at least six groups: minimally yellow for her, maximally yellow for her, minimally green for her, maximally green for her, minimally blue for her, maximally blue for her, minimally red for her, and maximally red for her. Miriam perceives the lemon as pure yellow in that she perceives it as belonging to a group including objects that are maximally yellow for her, minimally blue for her, minimally red for her, and minimally green for her. In contrast, she perceives a tangerine as a mixture of red and yellow in that she perceives sufficient similarities to maximally yellow objects and sufficient similarities to maximally red objects. In this way, the six aforementioned groups function like orientation points that she uses to locate all other objects, just as Ruth might use as orientation points (a) the rooms in which the mercury drops to its lowest possible level, and (b) the rooms in which the mercury rises to the top of the tube. An advantage of thinking of these groups as the orientation points is that it nicely explains our perceptions of colors as pure or as mixed, e.g., as pure yellow or as greenish yellow. I also find it introspectively plausible that our appreciation of an object's color involves an appreciation for how it relates to objects at the extremes.

Which objects are in these groups? A natural suggestion is that an object is minimally yellow for Miriam if she hasn't perceived it as yellower than any other object, and maximally yellow for Miriam if she hasn't perceived any other objects as yellower. Likewise for green, blue, and red. A potential disadvantage with this suggestion is that an object that is minimally yellow for her at one time might not be minimally yellow for her at a later time, because she might later perceive it as yellower than a new object. I'm not sure whether this is a disadvantage; I go back and forth. Another suggestion is that an object is minimally yellow for Miriam if she *can't* perceive it as yellower than anything else. Even if Miriam has never actually perceived an object that is minimally yellow for her, she might have an appreciation for how such an object would affect her, and therefore be able to make comparisons involving it. Analogously, even if Ruth has never seen the mercury in her thermometer drop to the very bottom, she might have an appreciation for how such a room would affect her thermometer, and therefore be able to make comparisons involving it. For example, she might say that her current room is close to that extreme, or halfway between that extreme and the opposite extreme. She might even compare all rooms indirectly, by first estimating each room's relations to the extremes. For example, by estimating one room as close to the extreme, and another room as halfway between the two extremes, she might rank the first room above the second. For our purposes, it won't ultimately be important which of these suggestions a structuralist adopts.

It also won't ultimately be important that I've correctly identified the six groups that function as orientation points, that these are the only groups that function as orientation points, or even that we always use the same orientation points. It's enough that we rely on such comparisons. For example, the structuralist might

instead insist that Miriam retains information about objects in all the groups, perhaps organizing them into a kind of color space, where the points are occupied by groups of objects rather than properties. In that case, there would be far more than six orientation points.

I think it's helpful to keep in mind that, like all the comparisons we've considered, these comparisons presumably result from processes hardwired into our brains. We don't *choose* to group an object together with other objects, any more than we choose to perceive the lemon as yellower than nearby objects and recently viewed objects.

Let's return to our original question. Miriam perceives the lemon as yellower than other objects. Which objects? We answered: objects she's currently perceiving, recently perceived, and perceived in the more distant past.

Importantly, I don't think she perceives relations to objects in all three categories with the same degree of specificity. Instead, I think she perceives the most specific relations to objects she's currently perceiving, less specific relations to objects she recently perceived, and the least specific relations to objects she perceived in the more distant past. She might perceive the lemon as exactly the same color as another lemon she's currently perceiving, but only roughly the same color as objects she perceived less recently. This comports with introspection. A perception of a monochromatic surface by itself is often less informative than a perception of that surface together with other surfaces. We're often unsure whether to report that such a surface is mustard yellow or straw yellow, even if we have had a lot of exposure to such objects. This might explain why simultaneous comparisons are an important part of paint selection. We perceive more specific colors when we can perceive an object's similarities to and differences from other visible objects. This might also explain people's behavior when, as described in the introduction, they're asked to turn a knob until an image on a screen is pure yellow. They don't turn the knob in one direction and then suddenly stop. Instead, they turn it beyond their eventual stopping place, as if checking for an image that looks even less green; they're relying on comparisons between images on both sides to select the pure yellow image. Finally, this might explain why, as Kalderon (2008, p.941) points out, our perception of an object's color can become more specific after we view it against additional backgrounds and under different illuminants. According to structuralists, these changes allow us to perceive additional relations.

I also don't think she always perceives objects within the same category with the same degree of specificity. The degree of specificity seems to depend on attention and spatial proximity, among other factors. I also suspect there's widespread interpersonal variation. As the result of innate differences, learning, or both, some people have a richer appreciation of how objects compare to other objects with respect to color, taste, and texture.

Finally, it's possible to *misperceive* any of these relations. Miriam can *misperceive* the lemon as yellower than an object that's nearby, an object that she recently perceived, or objects that she perceived in the more distant past. These misperceptions can result from external factors, such as a tinted illuminant. But they can also

result from internal factors, such as distortions in the processes underlying her comparisons to recently perceived objects.

6.4 *Matching*

Just as Ruth can't tell us anything about what it is for a room to be hotter than another room, Miriam can't tell us anything about what it is for an object to be yellower than another object. But we, *qua* theorists, can. Let's start with a preliminary sketch of what it takes for Miriam's perceptions to match, and thus for her to perceive, actual similarities and differences between objects.

By the time Miriam perceives the lemon, she has perceived millions of objects, many of them as yellower than other objects. Concatenate all those comparisons, adding information about the spatiotemporal locations of the relevant objects. Where ℓ_1, ℓ_2 , etc., are names for spatiotemporal locations, the list might start:

The lemon at ℓ_1 is yellower than the grapefruit at ℓ_2 , the grapefruit at ℓ_2 is yellower than the lime at ℓ_3, \dots

Next, consider all the relations between these objects. For example:

The lemon at ℓ_1 is more acidic than the grapefruit at ℓ_2 , the grapefruit at ℓ_2 is **not** more acidic than the lime at ℓ_3, \dots

The lemon at ℓ_1 is **not** denser than the grapefruit at ℓ_2 , the grapefruit at ℓ_2 is denser than the lime at ℓ_3, \dots

The lemon at ℓ_1 is smaller than the grapefruit at ℓ_2 , the grapefruit at ℓ_2 is **not** smaller than the lime at ℓ_3, \dots

Miriam's perceptions don't match any of these relations. Thus, when she perceives a lemon as yellower than a grapefruit, she's not perceiving the lemon as smaller, denser, or more acidic. But her perceptions might match another relation. That is, there might be a relation \mathcal{Y} such that:

The lemon at ℓ_1 bears \mathcal{Y} to the grapefruit at ℓ_2 , the grapefruit at ℓ_2 bears \mathcal{Y} to the lime at ℓ_3, \dots

In that case, when Miriam perceives the lemon as yellower than the grapefruit, she's perceiving it as \mathcal{Y} er, though not in a way that reveals the nature of that relation. Likewise, her perceptions of objects as redder, greener, bluer, etc., might match relations $\mathcal{R}, \mathcal{G}, \mathcal{B}$, etc. In the next subsection, I'll suggest that $\mathcal{Y}, \mathcal{R}, \mathcal{G}, \mathcal{B}$, etc., are differences in how objects reflect light. But for now, let's keep the discussion more general. To simplify, let's also set aside that Miriam perceives objects as yellower than each other to various degrees.

Miriam's perceptions are unlikely to *perfectly* match any relation. Perceptual illusions are one cause of mismatch. In Kitaoka's illusion, Miriam might perceive the left square as darker, but if she perceived the same squares in another context, she might perceive them as the same color. Memory failures are another cause of mismatch. Miriam's appreciation of how an object compares to past objects

might be mistaken, perhaps because she systematically misremembers objects as slightly yellower. Mismatch will also result from a certain kind of sorites series. When confronted with a series of pairwise indistinguishable objects, Miriam might perceive each pair as the same color, but the last and first objects as different colors. More generally, Miriam's perceptions will include a lot of mistakes and inconsistencies, making it unlikely that there will be a relation that perfectly matches her perceptions. Thus, we should be looking for the *best* match, not a perfect match.

Whether Miriam's perceptions match one relation, rather than another, depends on contingent facts about the objects in her environment, and thus will vary across possible worlds. There's a possible world in which Miriam has indistinguishable perceptions, but the detectors in her eye are sensitive to other kinds of differences, such as differences in how objects reflect x-rays and gamma rays, or perhaps even differences in acidity, because in that world acidity is remotely detectable and the grapefruit-like objects are more acidic than the lemon-like objects. In such a world, Miriam's perceptions will match a different relation. In some possible worlds, Miriam's perceptions won't match any relation at all, because none of the potential matches are close enough. Consider a world in which the phenomenal characters of her perceptions are completely random. In such a world, her perceptions won't match any of the actual similarities and differences between objects, and thus she won't be perceiving any relation between objects, just as, according to some philosophers, a person whose experiences are qualitatively indistinguishable from your experiences, but who lives in a world without liquids ("Dry Earth"), isn't thinking about any particular kind of liquid (for relevant discussion, see Boghossian 1997 and Pryor 2007, p.184–186). Miriam's perceptions describe a certain kind of role, and in some worlds, a different relation plays that role, and in other worlds, no relation plays that role.

Miriam's perceptions also won't always match one relation better than all the others. There's a possible world where size and acidity are nomically connected. In that world, her perceptions might equally match *is larger than* and *is more acidic than*. If she perceives a lemon-like object as yellower than a grapefruit-like object, I think we should say that it's indeterminate whether she's perceiving the lemon-like object as larger or as more acidic than the grapefruit-like object, because both relations equally match her perceptions. I think the situation in the actual world is similar, though less extreme. For reasons I'll introduce later, I think that in the actual world Miriam's perceptions match more than one difference in how objects reflect light. Thus, in the actual world, it's indeterminate which difference she's perceiving when she perceives the lemon as yellower than a lime.⁸ There are many ways of modeling this indeterminacy. We could, for example, use a supervaluational model (e.g., McLaughlin and McGee 1995). Nonetheless, because the structuralist response to the challenge of perceptual variation doesn't depend on any particular way of modeling indeterminacy, we don't need to pursue this issue here. For our purposes, we also don't need to worry about the best way to exclude unnatural, gerrymandered relations between objects. These are topics for another occasion.

Miriam's perceptions will *gradually* match some relations better than others. I think it's again helpful to consider a parallel between Ruth and Miriam. When

Ruth first started using ‘hotter’, it might have matched too many relations to have a specific environmental content. After the first three rooms, her comparisons might have equally matched differences in temperature, altitude, and volume. But after enough comparisons, her comparisons might have matched differences in temperature better than any of these other differences. I think Miriam’s situation has a similar structure. When Miriam first perceives an object as yellower than another, her perceptions won’t match any particular relation, so she won’t perceive any particular relation. But over time, her perceptions will match some relations better than others.

It might be helpful to think of Miriam’s perceptions as involving a symbol like *yellower*. When her perceptions first involve this symbol, she doesn’t perceive any particular relation. Her perceptions are like sentences that contain uninterpreted symbols. But over time, *yellower* starts to represent a more and more specific relation, and that’s what she perceives. Thus, when structuralists describe Miriam as perceiving the lemon as yellower than the lime, they’re not saying that she’s perceiving the same relation that she perceived in the past, or would perceive in another possible world. They’re saying that she’s perceiving whatever relation is currently best matched by her symbol *yellower*.

I think it helps to keep in mind that Miriam’s perceptions are the result of hardwired processes that are insensitive to whether symbols such as *yellower* are environmentally contentful. Thanks to these processes, Miriam automatically perceives lemons as yellower than other objects, even if her perceptions have yet to acquire a specific environmental content. At least in principle, she can sort yellow from green objects the moment she first opens her eyes; she can perceive that there is a difference without perceiving any specific difference. There’s also no limit to the number or complexity of these symbols, and thus no limit to the dimensions along which objects could be sorted. It’s just that she won’t perceive any particular difference between those objects until she’s sorted enough objects. Structuralists can insist that perception is a gradual achievement without thinking of our initial perceptions as too simple to guide complex sorting behavior. Analogously, we’re hardwired to sort faces, and there might be innate symbols like *face-29*. But these symbols don’t have innate meaning. They become contentful only after they’ve been activated by actual faces, because until that happens you can’t use these symbols to recognize any particular faces (“her again”).

So far, I haven’t said very much about Miriam’s phenomenal characters. I think that, as Miriam perceives increasingly specific relations, the environmental content of her phenomenal characters become more specific as well. In particular, at each moment the environmental content of Miriam’s total phenomenal character includes all the comparisons it elicits. The environmental contents of individual phenomenal characters, such as phenomenal-pure-yellow, are derivative. There’s *a lot* more to say about this issue. However, for the purposes of this paper, I just want to emphasize that the environmental contents of Miriam’s phenomenal characters gradually become more specific, along with the relations she’s perceiving.

This approach to perception is indebted to Ramsey’s, Carnap’s, and Lewis’s approach to theoretical identifications (see Lewis 1970). But there are also differences.


To start, the concatenation of all of Miriam's perceptions isn't a *theory*, at least in the sense that chemistry, economics, and folk psychology provide us with theories. It isn't predicative, explanatory, or universal. Relatedly, the concatenation of all of Miriam's perceptions doesn't include any *causal* or *functional* relations. It just includes relations such as *yellowier*, *greener*, *the same as*, *darker than*, etc. Finally, the concatenation of all of Miriam's perceptions grows over time, whereas the theories Ramsey, Carnap, and Lewis have in mind are much more stable.⁹

Structuralism has a number of advantages. Of particular interest, it doesn't push us to designate any contexts as "normal" or "ideal," because *all* Miriam's past perceptions are taken into account. It also doesn't push us to designate any observers as "normal" or "ideal," because *only Miriam's* past perceptions are taken into account. This is one of the reasons why structuralism doesn't commit us to ignorance about which contexts and observers are relevant. Moreover, by focusing on the past perceptions of each observer, structuralists are better able to explain how different observers, such as Miriam and Aaron, can accurately perceive the same object. I'll return to this point in subsection 7.

6.5 Light

Like structuralism itself, the approach we just considered is officially neutral about the metaphysics of color. If you think that objects have primitive, non-physical color properties, you'll think that Miriam perceives similarities and differences involving those properties, because they're the best match. For reasons that go beyond the scope of this paper, I don't think that objects have primitive, non-physical properties, and therefore I don't think Miriam perceives similarities and differences involving them. Instead, I think she perceives similarities and differences in how objects reflect light (i.e., wavelengths roughly between 400nm and 700nm). I'll therefore develop structuralism along these lines. But nothing will ultimately depend on this metaphysical assumption, at least with respect to the challenge of perceptual variation.¹⁰

Some scientific background might be helpful. Objects reflect different proportions of light at different wavelengths. Here are two possible ways that objects might reflect light:



At least in principle, objects can reflect light in infinitely many different ways, because they can reflect different proportions of light at any of the infinitely many wavelengths between 400nm and 700nm. As a result, there are a great many ways in which we could rank objects by how they reflect light. For example, we could rank

them by the proportion of light they reflect at 450nm. Alternatively, we could rank them by the proportion of light they reflect at both 450nm and 550nm, where more weight is placed on the proportion of light they reflect at 450nm. Or, more relevantly for our purposes, we could rank them by the proportion of light they reflect at *each* of the wavelengths between 400nm and 700nm, where different weights are placed on different wavelengths.

In the actual world Miriam's perceptions will eventually match a relation that ranks objects in a certain way. The amount of weight placed on each wavelength will be a function of the sensitivities of the relevant kinds of detectors in her eyes and the way her brain integrates information from those detectors. When she perceives the lemon as yellower than another object, she'll perceive the lemon as bearing that relation, though not in a way that reveals anything substantive about its nature, including that it involves differences in how objects reflect light.

As noted above, it's unlikely that Miriam's perception will match one relation better than all others. In part, that's because there are *infinitely many* different ways of ranking objects according to how they reflect light, and Miriam has perceived only *finitely* many objects. Her perceptions might equally match two relations, even though one places slightly more weight on the proportion of light reflected at a certain narrow range of wavelengths. Thus, it will be indeterminate which relation she's perceiving. Nonetheless, according to many models of indeterminacy (e.g., McLaughlin and McGee 1995), her perceptions will often be accurate, because, for example, the lemon bears all those relations to the lime.

6.6 Color

According to structuralists, Miriam perceives the lemon as a certain color because she perceives it as yellower, etc., than other objects. Which color? If Miriam's perceptions best match differences in how objects reflect light, then colors are ways of reflecting light. Thus, she's perceiving the lemon as reflecting light in a certain way. In what sense does Miriam perceive the lemon as having that color *because* she perceives it as yellower, etc., than other objects? As I think structuralism should be developed, she perceives it in roughly the sense that the definite description 'the tallest man in the room is bald' describes Boaz, because he's taller than the other men in the room.¹¹

A mathematical model will help clarify these claims, including the sense in which our perception of the lemon as yellower than the lime is like a description of its color. Suppose a , b , and c are points in the unit interval, $[0,1]$, and I merely represent a as .2 to the left of b . I am thereby describing a as in $[0,.8]$ and b as in $[.2, 1]$. Suppose I then also represent c as .2 to the right of b . I am thereby describing a as in $[0,.5]$, b as in $[.2, .7]$, and c as in $[.5, 1]$. The more relations I represent, and the more specific those relations, the more specific the intervals I describe a as occupying. Once I've represented enough relations among points, perhaps even millions of relations among millions of points, representing additional points won't significantly increase the specificity of the intervals.

I think that color perception is similar: Miriam perceives the lemon's color in the same sense that I described a 's location. The details, however, are far more

complicated, at least if colors are ways of reflecting light. Perhaps most significantly, our model involved a line segment, $[0,1]$. In contrast, the space of reflectances is $[0,1]^\infty$, because there is a dimension for each wavelength between 400nm and 700nm. As a result, whereas in the mathematical model we represent a one-dimensional interval containing a , in color perception Miriam perceives an infinite-dimensional region that contains the lemon's way of reflecting light. I will describe those regions as *determinables*, because I think that the region itself is a color, rather than a mere disjunction of colors. But I won't assume anything controversial about the nature of determinables; for my purposes, it's just a convenient way of talking. As in our model, the more relations Miriam perceives involving the lemon, and the more specific those relations, the more specific the color she perceives.

Of all the views we considered in earlier sections, this view is closest in spirit to Jackson, Pargetter, and McLaughlin's, because the colors are whatever play a certain kind of role. But it's importantly different, because the roles involve relations that are non-causal, become more specific over time, and relate objects to each other, rather than to a perceiver.

There's also a similarity to Fregean views of perceptual representation. On the view I'm developing, there are two levels of representation. On the first level, there are representations of color relations. On the second level, there are representations of monadic color properties. Because the first level determines the second, it might be tempting to think of the first level as a Fregean sense. But I don't think that's a helpful way of thinking about structuralism, because the color relations and color properties are represented *together*, as part of a unified scene. Traditionally, however, we're not said to perceive Fregean senses alongside their referents.

6.7 Structural Both-ism

Finally, the payoff: Structuralism gives us an attractive way of defending the suggestion that both Miriam and Aaron are accurately perceiving the lemon. The lemon is at the very top of the ranking that best matches Miriam's perceptions. In that ranking, higher objects bear \mathcal{Y} to lower objects. Because of differences between Miriam's and Aaron's eyes and brain, the lemon is not at the very top of the ranking that best matches Aaron's perceptions. In this other ranking, higher objects bear \mathcal{Y}' to lower objects. As a result, when Miriam perceives the lemon as yellower than another object, she perceives it as \mathcal{Y} er than the other object, and when Aaron perceives another object as yellower than the lemon, he perceives it as \mathcal{Y}' er than the lemon. Both of these perceptions might be accurate. Speaking loosely, the lemon might reflect a greater proportion of light at wavelengths that matter to the \mathcal{Y} -ranking while the other object reflects a greater proportion of light at wavelengths that matter to the \mathcal{Y}' -ranking. Thus, both Miriam and Aaron might be accurately perceiving relations between the lemon and other objects.

Miriam and Aaron might also be accurately perceiving the color of the lemon. I think that our mathematical model is again helpful. As before, suppose I describe a as in $[0,.5]$ by describing its *distances* to b and c . You might describe a as in $[0,.5]$ using a different relation. Whereas the distance between a and b is the value of the function $|a - b|$, let the *schmistance* between a and b be the value of the function

$1 - |a - b|$. You might describe a as in $[0, .5]$ by describing the schmistances between a , b , and c . In particular, you might describe a as in $[0, .5]$ by describing it to the left of b with a schmistance of .8, and b to the left of c with a schmistance of .7. Thus, if a is actually located in $[0, .5]$, we both might have accurately described its location, even though we're using different relations. It also doesn't matter if we're describing relations to different objects. You can describe a as in $[0, .5]$ by describing it to the left of d with a schmistance of .9, and d to the left of e with a schmistance of .6. Finally, it might not even matter if we're describing different intervals. We might be describing intersecting regions such that a is in the intersection. Even if I describe a as in $[0, .5]$ and you describe it as in $[.3, .6]$, we're both accurately describing a 's location if it is located within the intersection.

Returning to the color case, Miriam and Aaron are perceiving different relations. In particular, they are perceiving \mathcal{Y} and \mathcal{Y}' , respectively. They are also perceiving relations to different objects. For example, Miriam is perceiving relations to objects that are maximally pure yellow for her, and Aaron is perceiving relations that are maximally pure yellow for him. Nonetheless, they might both be accurately perceiving the lemon's way of reflecting light, and therefore both be accurately perceiving its color. One possibility is that they're perceiving intersecting regions that both contain the lemon's way of reflecting light. In other words, they might be perceiving determinables that include the lemon's way of reflecting light but don't include all the same ways of reflecting light. The details are far more complicated than in our mathematical model, because the determinables are infinite dimensional regions, and also because there might not be a higher-order function that maps \mathcal{Y} to \mathcal{Y}' or vice versa. But the basic point is the same. Thus, structuralism gives us a new way of developing both-ism.

Importantly, structuralism doesn't imply that *whenever* there is interpersonal variation, both perceivers are accurately perceiving the object's color relations and individual color. Aaron's perceptions might have instead matched a relation \mathcal{Y}'' so that Aaron perceives the lemon as \mathcal{Y}'' er, even though it's not. More abstractly, because Aaron perceives the relation that *best* matches his perceptions, rather than the relation that *perfectly* matches his perceptions, in some cases he'll misperceive. Thus, in cases of interpersonal variation, sometimes only one person's perception is inaccurate, and sometimes neither person's perception is accurate.

It can also be indeterminate whether Aaron is accurately perceiving the lemon, because it can be indeterminate whether his perceptions match \mathcal{Y}' or \mathcal{Y}'' , and the lemon might be \mathcal{Y}' er but not \mathcal{Y}'' er than the lime. Thus, in cases of interpersonal variation, it's sometimes indeterminate whether a certain person's perception is accurate, and it is also sometimes indeterminate whether anyone's perception is accurate.

The structuralist's response to the challenge of perceptual variation has a number of advantages over the other responses we considered. First, unlike those who respond "neither," the structuralist can explain why our color perceptions are sometimes accurate, thereby preserving our natural ways of thinking and talking, at least to some extent.

Second, unlike those who respond "one," the structuralist response doesn't lead to unacceptable ignorance. Provided we know all the relevant facts about Miriam's

past perceptions, including details about the objects she perceived, we can figure out which relations she's perceiving, and thus whether her perceptions are accurate. Of course, we rarely if ever know all the relevant facts. But that's acceptable ignorance, because our ignorance has the right kind of causal explanation. We might be ignorant because most of the information about Miriam's past perceptions has been lost (as with our ignorance of Socrates's exact height at the time of his death) or because she has too many perceptions each minute for us to keep track (as with our ignorance of the weight of all the chocolate in the world). Whatever the explanation, this is an acceptable kind of ignorance.

Third, unlike atomists who respond "both," the structuralist can explain perceptual error. As noted above, the best match is almost always imperfect. As a result, there are almost always cases of misperception, even in contexts that are normal by any reasonable standard. Thus, even if Miriam and Aaron are both accurately perceiving the color of the lemon, there are other objects whose colors they misperceive, even in normal contexts, and even though they're both normal perceivers. In this way, structuralists can explain perceptual inaccuracy.

Kitaoka's lightness illusion is again helpful. In that illusion, you perceive the left square as darker than the right square. You also perceive those squares as standing in different relations to nearby objects (e.g., your hand), objects you perceived recently, and objects you perceived in the more distant past. As a result, your perceptions of each square's relations describe a different color. Nonetheless, the squares are intrinsically alike. Thus, you're misperceiving at least one of the squares. I think this is the right result. The most natural thing to say about Kitaoka's illusion is that, because you're misperceiving the left square as darker, you're misperceiving at least one of the squares. This is the kind of thing we're inclined say in other domains: if you perceive one person as shorter than another, and they're really the same height, the natural conclusion is that you're accurately perceiving at most one of their heights. Of course, it isn't immediately obvious which square you're misperceiving. To determine that, we'd have to know all of the relevant facts about your past perceptions, including details about the objects you perceived. But there's still a knowable fact of the matter, even if we don't know it.

Finally, unlike those who respond "indeterminate," the structuralist does not need to explain why there is so much indeterminacy that it's impossible for people to accurately perceive colors. As an added benefit, structuralists give us a nice, bottom-up explanation of why there's indeterminacy in what we perceive.

For these reasons, I think that structuralism gives us the most promising response to the challenge of perceptual variation. As I noted in the introduction, I also think it gives us promising solutions to other puzzles (see my 2012, 2013). But perhaps most importantly, I think it gives us a coherent and plausible way of thinking about perception.

6.8 Features

Structuralism has a number of interesting features. For example, it creates a fundamental division between perception and thought. Suppose a soft matter physicist tells you that the dielectric constant of paper is 3.85. Even if this is the first time

you've heard the phrase 'dielectric constant' and you don't have any beliefs about which materials have lower or higher dielectric constants, you can still believe that the dielectric constant of paper is 3.85 on the basis of this testimony. Thus, in thought you can represent an object's quality without representing any relations to other objects. According to structuralists, perception is different, at least with respect to secondary qualities. You can perceive the lemon as yellow only by representing it as yellower than something else.

Structuralism also gives us a new way of reconciling what we perceive with what science tells us. According to our best physical theories, we live in a world made up of photos, bosons, and electrons. But according to our perceptions, we live in a world containing yellow lemons and sweet chocolates. Many have wondered: How can these be one and the same world? A traditional, atomistic approach is to say that we perceive the causes of our perceptions. More precisely, we perceive whatever causes those perceptions in a certain kind of perceiver. Given that our best physical theories completely describe the causes of evaporation and photosynthesis, it's natural to think that they also completely describe the causes of perceptions in the relevant kind of perceiver. In that case, what we perceive must be included in our best scientific description of the world. Despite the initial plausibility of this atomistic approach', I think that, like all atomistic proposals, it leaves us without a satisfying response to the challenge of perceptual variation.

Structuralism gives us an alternative. Structuralists encourage us to think of perception as a map, where over time more and more objects are placed further or closer along a number of dimensions, until there's enough detail that it becomes an imperfect map of properties and relations in the external world. Assuming the completeness of our best scientific descriptions of objects in the external world, this allows us to reconcile what we perceive with what science tells us, because our perceptions will best match relations included in our scientific descriptions. Structuralists thus deny that we perceive something outside us merely because it *causes* the right kind of change inside us. They instead claim that perception requires a *structural match*, and causal relations aren't sufficient for structural matches. Moreover, while causal relations might still be necessary, their only role is to explain why we perceive some objects rather than others, e.g., this lemon, rather than one of the other lemons scattered across the globe. Structuralism thereby gives us a more promising way of reconciling what we perceive with what science tells us.

Structuralism is also potentially naturalistic. When explaining why Miriam perceives the lemon as pure yellow, many appeal to causation or natural selection, because these relations are uncontroversially natural, in that they play important roles in our best scientific theories. Structuralists integrate perceptual intentionality into the natural world in another way. Structuralists appeal to matching, a structural relation that doesn't require anything unnatural or extra-scientific. Much like *is rounder than* and *has the same number of parts as*, it's a relation between wholly natural things. Structuralists also appeal to mental relations, like retention and grouping, that seem functional in nature, and thus naturalistic. Of course, there's much more to say about both issues. But, for now, I just want to point out that

structuralism is potentially naturalistic, even though it has little in common with what Dretske (1981), Fodor (1987), and Millikan (1984) propose.

Structuralism is also a new view about the role of phenomenal characters in perception. According to structuralists, phenomenal characters are like meaningless symbols that become meaningful over time, as the result of comparisons. Is structuralism therefore inconsistent with representationalism, the view that what we perceive supervenes on our phenomenal characters? As I think structuralism should be developed, what Miriam perceives changes over time, even though her phenomenal characters remain the same, e.g., some of her perceptions are still phenomenal-pure-yellow. So developed, structuralism is inconsistent with representationalism. But there are other ways of developing structuralism, some of which are consistent with at least certain versions of representationalism. Perhaps as she perceives more specific properties, her perceptions acquire new phenomenal characters, even though she's unable to detect that change through introspection. In that case, what she perceives might still supervene on her phenomenal characters. Alternatively, structuralism might be developed so that Miriam's perception also represents what Shoemaker (1994) calls "appearance properties" or what Chalmers (2006) calls "Edenic colors." In that case, structuralism would be consistent with some versions of representationalism, in that some of the properties we perceive would supervene on our phenomenal characters.

Finally, structuralism generalizes beyond color perception, at least to the other secondary qualities. In each case, there's a distinctive kind of phenomenology that allows us to perceive similarities and differences between objects. In the case of taste, we might perceive objects as saltier or sweeter, and thereby eventually perceive similarities and differences between the chemicals in what we're eating.

Structuralism might also generalize to our perception of space. Even if there isn't a distinctive kind of spatial phenomenology, our perceptions of spatial patterns (e.g., collinear) and of relative distances (e.g., twice our eye height) might eventually match certain relations in the physical geometry of our world.

Structuralism might even generalize to object recognition. As I suggested in Section 6, it might explain why I recognize a person *as Isaac*, and a lemon *as a lemon*.

But it can't explain everything. In particular, it can't explain why I'm perceiving *this lemon*, rather than another lemon. To explain that, I think we'll need to appeal to causal relations. But structuralism might be able to explain the rest, including why I accurately perceive it as a yellow lemon located roughly five meters away. Structuralism gives us a remarkably general framework for explaining our perceptions of the external world.

7. Conclusion

Philosophers have been working out the details of atomism for a long time. It will take structuralists a while to catch up. We'll have to leave many questions and objections unanswered, at least for now, including questions and objections involving unconscious perception, phenomenal inversion, representationalism, swampmen,

matching, gerrymandering, and the veridicality of memory. But I hope I've said enough to convince you that these questions and objections are worth our time.¹²

Notes

¹ The canonical example, attributed to Democritus, is that honey tastes sweet to those who are healthy, but bitter to those who are sick. See Sextus Empiricus 2000, Bk I, §101, §211–213; Lucretius 1924, Bk 4, ln 642–672, and Plutarch 1967, 1120E. Sextus discusses Democritus and Heraclitus. When ancient authors ask us to choose between color perceptions, they usually ask us to choose between the perceptions of observers in different contexts, such as perceptions of the pigeon's neck from different angles, the peacock's tail from the front and back, and the sea when it is rough and calm (see Ierodikonou 2015 for extensive citations). But there are exceptions. Among them: Lucretius 1924, Bk IV, ln 332–336 contrasts the color perceptions of a healthy person and someone with jaundice; Plutarch 1967, 1109C–E contrasts the perceptions of someone with normal eyes and someone with misshapen eyes; and Sextus Empiricus 2000, Bk I, §105 contrasts the perceptions of the young and old, claiming that the young perceive brighter colors.

² See Sextus Empiricus 2000, Bk I, §102–103; Bk II, §54.

³ See Byrne and Tye's (2006, p.252–253) reply to Pautz (2006, p.214–228). For more advantages to focusing on actual variation, see Morrison manuscript b, fn 4.

⁴ See Stroud 2000, p.173–176; Tye 2002, p.108; Byrne and Hilbert 2003, p.16–17; 2004, p.37–39, 2007, p.88, Byrne and Tye 2006, p.252; and Allen 2016, p.58–65. I don't think that one-ism's critics (e.g., Cohen 2009, p.45–64) have adequately explained what's wrong with this response. For my explanation, see Morrison manuscript b.

⁵ Another alternative is that there is a *contradiction* between how we perceive the colors of the squares and how we perceive their relation. In particular, there is a color *c* such that you perceive the left square as *c* and the right square as *c*, but nonetheless perceive the left square as darker. I address this in Morrison (manuscript a).

⁶ Other alternatives: Hilbert and Kalderon (2000, Sect 5) explain facts like (1) by appealing to dispositions involving perceptions of many kinds, and then explain facts like (2) by appealing to facts like (1) (see in particular p.200–201). Rosenthal (2005) also explains facts like (1) by appealing to dispositions involving perceptions of many kinds, but doesn't offer an explanation of facts like (2). Matthen (2010) explains facts like (1) by appealing to the dispositions responsible for conditioning, in particular the spread of learned expectations from one color experience to another, and then explains facts like (2) by appealing to facts like (1). While all these authors appeal to "color spaces," this is just a first step toward explaining facts like (1), and doesn't yet involve an explanation of facts like (2). Their explanations have the following order: dispositions involving perceptions of many kinds → color space → facts like (1) → facts like (2). Clark (2000, Ch 6) denies that there are facts like (1) (p.233–234), because he's dissatisfied with atomistic explanations of such facts (p.208–228). He instead claims our perceptions are exhausted by facts like (2) (though he sometimes seems to say we also perceive non-physical, uninstantiated colors—see p.245). He also claims that perceptions of color similarities are usually illusory (p.225). Structuralists, in contrast, explain facts like (1) and argue that we often accurately perceive colors and color similarities. For other differences between structuralism and the alternatives from Hilbert and Kalderon (2000) and Matthen (2010), see footnote 9.

⁷ Cohen 2008 argues that Miriam represents two kinds of similarities and differences: those that don't discount for contextual factors, and those that do. I disagree, but it doesn't matter, because structuralists can give parallel accounts of these two kinds of relations.

⁸ There might be other sources of indeterminacy. If we place more weight on her more recent perceptions, to allow her to accurately perceive relations despite recent changes in her eyes (e.g., cataracts), it might be indeterminate how we should place those weights. Also, due to indeterminacy in the notion of matching, there might not always be a determinate fact about which relation best matches Miriam's perceptions.

⁹ There's also a resemblance between structuralism and the views of Hilbert and Kalderon 2000, p.196–202, and Matthen 2010, p.77–88, in that all these views prioritize comparisons (see also Kalderon

2007, p.590–595, 2011, p.242–245). But there are many important differences. First, structuralists appeal to Miriam's *past perceptions*, while Hilbert, Kalderon, and Matthen appeal to her *innate dispositions*. In particular, Hilbert and Kalderon focus on her innate dispositions to compare objects, and Matthen focuses on whatever innate dispositions are responsible for the spread of learned expectations from one color to another. In either case, these dispositions manifest differently in different contexts, and thus Hilbert, Kalderon, and Matthen need to justify their focus on a particular context, or assortment of contexts, at least if they want to avoid the kind of ignorance that led us to reject one-ism. By focusing on all of Miriam's past perceptions, structuralists don't have this burden, because they don't focus on any particular context. Second, structuralists appeal to a *structural match*, which seems to be a more specific kind of mind-world relation than what Hilbert and Kalderon call "selection." Hilbert and Kalderon don't say much about how our innate dispositions "select" specific relations between objects. Third, structuralists claim that Miriam perceives the lemon as yellow because she represents the lemon as yellower than *other objects*, while Hilbert, Kalderon, and Matthen claim that Miriam perceives the lemon as yellow because she represents *high-order* relations between the lemon's color and *other properties*, in particular the other properties in the space of color properties. Fourth, as noted earlier (fn 6), Hilbert, Kalderon, and Matthen accept the second constitutive claim of atomism, in that they claim that Miriam perceives the lemon as yellower than the lime because she perceives the lemon as pure yellow and the lime as pure green. They just reject the first constitutive claim of atomism, that Miriam perceives the lemon as pure yellow because of a relation involving perceptions of that kind, rather than a relation involving perceptions of many kinds.

¹⁰ Notably, structuralists aren't the only philosophers who can accept this kind of account, at least in principle. Atomists might ask which *property* best matches Miriam's perception of objects *as yellow*, rather than which *relation* best matches Miriam's perception of objects *as yellower*. If atomists take this approach, I think they'll end up with a view like Jackson, Pargetter, and McLaughlin's. In particular, if to perceive an object as yellow is to have a perception with a phenomenal-yellow character, they'll need to restrict their focus to a given kind of context, because virtually any object can cause a phenomenal-yellow perception, given the right lighting, distance, background, etc. Thus, like Jackson, Pargetter, and McLaughlin, they'll think that the colors are what satisfy descriptions like 'the cause of phenomenal-yellow perceptions in Miriam in a given context'. I criticize such views above, in Section 5, and also in Morrison (manuscript a). That being said, this might be an opportunity for philosophical innovation. The challenge is to come up with another reductive, non-intentional understanding of what it is to perceive an object *as yellow*, besides to have a perception with a phenomenal-yellow character.

¹¹ This isn't the only option available to structuralists. If she's perceiving the lemon as yellower than a lime, and if her perceptions match a relation \mathcal{Y} , another option is that she's perceiving the lemon as instantiating *is \mathcal{Y} er than the lime*. If she were also perceiving the lemon as yellower than a tangerine, she'd be perceiving the lemon as instantiating *is \mathcal{Y} er than the lime and is \mathcal{Y} er than the tangerine*. This is in the spirit of Cohen's proposal, because the relevant relations are included in the color.

There are several senses in which she might represent this property *because* she perceives the lemon as yellower than the lime. Perhaps there's no difference between perceiving the lemon as instantiating this property and representing the lemon as \mathcal{Y} er than the lime. Likewise, one might think that there's no difference between believing *that someone is kissing Boaz* and believing *that Boaz is being kissed* even though the first belief is about a relation (*is kissing*) and the second belief is about a property (*is being kissed*). One might think that these are just different ways of reporting the same belief. Alternatively, if there is a difference between perceiving the lemon as \mathcal{Y} er than the lime and perceiving the lemon as instantiating *is \mathcal{Y} er than the lime*, the first perception might cause the second. Perhaps we perceive many such relations but only a few colors. We might perceive an object's individual color just when we attend to it. If we're thinking of the mind in computational terms, the perceptual system might introduce a temporary symbol for the property of the relevant object (THAT COLOR), which it then stores or discards.

If we accept this option, we'll give a slightly different explanation of how both Miriam and Aaron can be accurately perceiving the lemon. We'll say that Miriam is perceiving a property like *is \mathcal{Y} er than this other lemon*, and Aaron is perceiving a property like *is just as \mathcal{Y}' as this other lemon*. More generally, we'll say they are perceiving properties that are derived from the relations they're perceiving. Because they both might be accurately perceiving the lemon's relations, they both might be accurately

perceiving the lemon's color. I explore this kind of view in more detail in my 2013, including its account of color constancy and illusion. I've since become more attracted to the view developed above, because it allows us to perceive the same color on different occasions, even though we're perceiving different relations to different objects, which I think better approximates our natural ways of thinking and talking.

¹² A talented and generous group of people helped this paper develop from its first draft in September of 2013 to its final draft in January of 2018. With apologies to those that I've forgotten, thanks to David J. Barnett, Alex Byrne, Justin Clarke-Doane, Jonathan Cohen, Shamik Dasgupta, Sinan Dogramaci, Cian Dorr, Matthew Fulkerson, Adam Pautz, and Susanna Schellenberg for their comments along the way. Most significantly, they convinced me to expand two sections (sections 3 and 5) into separate papers (manuscript a and manuscript b). Thanks also to audiences at Barnard, Yale, and NYU for their terrific questions.

References

- Allen, Keith (2016). *A Naive Realist Theory of Colour*. Oxford University Press UK.
- Armstrong, D. M. (1978). *Nominalism and Realism: Volume 1: Universals and Scientific Realism*. Cambridge University Press.
- Bell, J. S. (2004). *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*. Cambridge University Press. First published 1987.
- Block, Ned (1999). Sexism, ageism, racism, and the nature of consciousness. *Philosophical Topics*, 26(1): 39–70.
- Boghossian, Paul A. (1997). What the externalist can know a priori. *Proceedings of the Aristotelian Society*, 97: 161–175.
- Byrne, Alex and Hilbert, David R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, 26(1): 3–21.
- Byrne, Alex and Hilbert, David R. (2004). Hardin, tye, and color physicalism. *The Journal of Philosophy*, 101: 37–43.
- Byrne, Alex and Hilbert, David R. (2007). Truest blue. *Analysis*, 67(293): 87–92.
- Byrne, Alex and Tye, Michael (2006). Qualia ain't in the head. *Noûs* 40(2): 241–255.
- Chalmers, David (2006). Perception and the fall from Eden. In Gendler, Tamar Szabó and Hawthorne, John, editors, *Perceptual Experience*, pages 49–125. Oxford University Press.
- Clark, Austen (2000). *A Theory of Sentience*. Oxford University Press.
- Cohen, Jonathan (2004). Color properties and color ascriptions: A relationalist manifesto. *The Philosophical Review*, 113: 451–506.
- Cohen, Jonathan (2007). A relationalist's guide to error about color perception. *Noûs*, 41(2): 335–353.
- Cohen, Jonathan (2008). Colour constancy as counterfactual. *Australasian Journal of Philosophy*, 86(1): 61–92.
- Cohen, Jonathan (2009). *The Red and the Real: An Essay on Color Ontology*. Oxford University Press.
- Dasgupta, Shamik (2016). Symmetry as an epistemic notion. *British Journal for the Philosophy of Science*, 67(3): 837–878.
- Dretske, Fred I. (1981). *Knowledge and the Flow of Information*. The MIT Press, Cambridge, MA.
- Empiricus, Sextus (2000). *Outlines of Scepticism*. Cambridge Texts in the History of Philosophy. Cambridge University Press.
- Fodor, Jerry A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. The MIT Press.
- Hilbert, David R. and Kalderon, Mark Eli (2000). Color and the inverted spectrum. In Davis, Steven, editor, *Color Perception: Philosophical, Psychological, Artistic, and Computational Perspectives*, volume 9 of *Vancouver Studies in Cognitive Science*, pages 187–214. Oxford University Press.
- Husserl, Edmund (1980). *On the Phenomenology of the Consciousness of Internal Time*. Kluwer Academic Publishers. Lectures given between 1893 and 1917.
- Ierodiakonou, Katerina (2015). Hellenistic philosophers on the phenomenon of changing colours. In Holmes, Brooke and Fischer, Klaus-Dietrich, editors, *The Frontiers of Ancient Science: Essays in Honor of Heinrich von Staden*, pages 227–250. De Gruyter.

- Ismael, Jenann and van Fraassen, Bas C. (2003). Symmetry as a guide to superfluous theoretical structure. In Brading, Katherine and Castellani, Elena, editors, *Symmetries in Physics: Philosophical Reflections*, pages 371–392. Cambridge University Press.
- Jackson, Frank and Pargetter, Robert (1987). An objectivist's guide to subjectivism about color. *Revue Internationale de Philosophie*, 41: 127–141.
- James, William (1950). *The Principles of Psychology*. Dover. First published in 1890.
- Johnston, Mark (2004). The obscure object of hallucination. *Philosophical Studies*, 120(1–3): 113–183.
- Kalderon, Mark Eli (2007). Color pluralism. *The Philosophical Review*, 116: 563–601.
- Kalderon, Mark Eli (2008). Metamerism, constancy, and knowing which. *Mind*, 117(468): 935–971.
- Kalderon, Mark Eli (2011). The multiply qualitative. *Mind*, 120(478): 239–262.
- Kitaoka, Akiyoshi (2006). Luminance-gradient-dependent lightness illusion. <http://www.psy.ritsumei.ac.jp/~akitaoka/gilchrist2006mytalke.html#LGLI>.
- Kriegel, Uriah (2013). Two notions of mental representation. In Kriegel, Uriah, editor, *Current Controversies in Philosophy of Mind*, pages 161–179. Routledge.
- Kuehni, Rolf G. (2004). Variability in unique hue selection: A surprising phenomenon. *Color Research and Application*, 29(2): 158–162.
- Lewis, David (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13): 427–446.
- Lucretius, (1924). *On the Nature of Things*. Loeb Classical Library. Harvard University Press.
- Matthen, Mohan (2010). Color experience: A semantic theory. In Cohen, Jonathan and Matthen, Mohan, editors, *Color Ontology and Color Science*, pages 67–90. The MIT Press.
- Maudlin, Tim (2012). *Philosophy of Physics: Space and Time*. Princeton University Press.
- McLaughlin, Brian (2003). Colour, consciousness, and colour consciousness. In Smith, Quentin and Jovic, Aleksander, editors, *Consciousness: New Philosophical Perspectives*, pages 97–154. Oxford University Press.
- McLaughlin, Brian and McGee, Vann (1995). Distinctions without a difference. *The Southern Journal of Philosophy*, 33: 203–251.
- Millikan, Ruth (1984). *Language, Thought, and Other Biological Categories*. The MIT Press.
- Morrison, John (2012). Colour in a physical world: A problem due to visual noise. *Mind*, 121(482): 333–373.
- Morrison, John (2013). Anti-atomism about color representation. *Noûs*, 47(2): 94–122.
- Morrison, John (manuscript-a). Perceptual variation and relativism.
- Morrison, John (manuscript-b). Perceptual variation and ignorance.
- Neitz, Jay and Jacobs, Gerald H. (1986). Polymorphism of the long-wavelength cone in normal human colour vision. *Nature*, 323: 623–625.
- Pautz, Adam (2006). Sensory awareness is not a wide physical relation: An empirical argument against externalist intentionalism. *Noûs*, 40(2): 205–240.
- Pautz, Adam (2010a). A simple view of consciousness. In Koons, Robert C. and Bealer, George, editors, *The Waning of Materialism*, pages 25–66. Oxford University Press.
- Pautz, Adam (2010b). Why explain visual experience in terms of content? In Nanay, Bence, editor, *Perceiving the World*, pages 254–309. Oxford University Press.
- Peacocke, Christopher (1984). Colour concepts and colour experience. *Synthese*, 58: 365–381.
- Plutarch, (1967). *Moralia*, volume 14 of *Loeb Classical Library*. Harvard University Press.
- Pryor, James (2007). What's wrong with McKinsey-style reasoning? In Goldberg, Sanford C., editor, *Internalism and Externalism in Semantics and Epistemology*, pages 177–200. Oxford University Press.
- Rosenthal, David M. (1986). Two concepts of consciousness. *Philosophical Studies*, 49(3): 329–359.
- Rosenthal, David M. (2005). *Consciousness and Mind*. Clarendon Press.
- Scheffrin, B. and Werner, J. S. (1990). Loci of spectral unique hues throughout the lifespan. *Journal of the Optical Society of America A*, 7: 305–311.
- Shoemaker, Sydney (1994). Phenomenal character. *Noûs*, 28: 21–38.
- Smith, Vivienne C. and Pokorny, Joel (2003). Color matching and color discrimination. In Shevell, Steven K., editor, *The Science of Color*, pages 103–148. Elsevier.

- Stockman, Andrew and Sharp, Lindsey T. (1999). Cone spectral sensitivities and color matching. In Gegenfurtner, Karl R. and Sharpe, Lindsay T., editors, *Color Vision: From Genes to Perception*, pages 53–88. Cambridge University Press.
- Stroud, Barry G. (2000). *The Quest for Reality: Subjectivism and the Metaphysics of Colour*. Oxford University Press.
- Tye, Michael (2002). *Consciousness, Color, and Content*. The MIT Press.
- Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, 2nd edition.