

*The Perspectival Character of Perception**

Kevin J. Lande

University of California, Los Angeles

Abstract

We can perceive things, in many respects, as they really are. Nonetheless, our perception of the world is perspectival. You can correctly see a coin as circular from most angles. Yet the coin looks different when slanted than when head-on, and there is some respect in which the slanted coin looks similar to a head-on ellipse. Many hold that perception is perspectival because we perceive certain properties that correspond to the “looks” of things. I argue that this view is misguided. I consider the two standard versions of this view. What I call the *pluralist approach* fails to give a unified account of the perspectival character of perception, while what I call the *perspectival properties approach* violates central commitments of contemporary psychology. I propose instead that perception is perspectival because of the way perceptual states are structured from their parts.

*I am indebted to Tyler Burge and Gabriel Greenberg for their comments on multiple drafts. I also received helpful comments from Katie Elliot, Bill Kowalsky, Gabbrielle Johnson, Gabe Dupre, Ian Boon, Rene Jagnow, Susanna Schellenberg, and Ned Block. Thanks also to those involved in the 2016 IIFs-UNAM Philosophy Graduate Conference.

In perception, one can transcend the limitations imposed by one's perspective and see things, in many respects, as they really are. For example, a circular coin can reflect dramatically different patterns of light to your eyes depending on where you are viewing it from. Each of these patterns of light could have been reflected by infinitely many different types of surfaces of different shapes and sizes. On the face of it, these ever-changing and always ambiguous patterns of light at the eye carry little information about the shape and size of the coin. Yet normally you can see the coin as having a unique size and shape. The size and shape you see the coin as having normally remains the same from one viewpoint to the next—this is known as *size constancy* and *shape constancy*. And, normally, you can see the size and shape of the coin accurately. But while you can perceive the size and shape of the coin accurately from just about any perspective, your perception of the coin remains marked by your perspective on it. What accounts for the perspectival character of perception?

Two types of phenomena exemplify the perspectival character of perception. Take a circular coin and rotate it away from you in depth. In the first place, your perception of the coin changes in some respect as you rotate it. Loosely speaking, the coin has a different look when seen head-on than when seen at an angle even though we see it as circular in both cases. Call this sort of phenomenon *perspectival variance*.¹ Second, it is often said that there is some respect in which your perception of the slanted coin is similar to your perception of a head-on ellipse. Loosely speaking, even as the coin looks circular to you, there is some sense in which it also has an elliptical look. Call this sort of phenomenon *perspectival similarity*.²

1. This sort of phenomenon has also been called “perceptual relativity,” “situation-dependency,” “non-constancy,” and “inconstancy.”

2. While most people acknowledge that there is perspectival variance, some doubt that there is perspectival similarity. See, for example, Arthur D. Smith, *The Problem Of Perception* (Cambridge, MA: Harvard University Press, 2002); Walter Hopp, “No Such Look: Problems with the Dual Content Theory,” *Phenomenology and the Cognitive Sciences* 12, no. 4 (2013): 813–833. For the purposes of this paper, I will treat

Perspectival variance and similarity can be found throughout perception.³ Your perception of non-geometrical properties like color and lightness exhibits patterns of perspectival variance and similarity.⁴ For example, your perception of a white surface under shade differs from your perception of that white surface under sunlight (perspectival variance), and some hold that your perception of the white surface under shade is similar in some respect to your perception of a darker gray surface that is under sunlight (perspectival similarity), even if your perceptions of both are fully accurate.⁵ Analogous cases arise in hearing and touch. While these forms of perception differ in many ways, they all raise the same general puzzles. Given that we can accurately perceive the sizes, shapes, and colors of things under many different conditions of observation, why should our perceptions of these properties vary depending on those conditions? And why should there be any relevant similarity between our accurate perceptions of two different properties (circularity and ellipticality, for example) when each is presented in a different condition of observation (slanted and head-on, respectively)? If, in perception we can transcend the ever-changing flux of equivocal sensations and represent the more or less fixed properties of things in the world, then why does perception reflect that changing, equivocal flux, in the form of variance and similarity?

The standard explanations of the perspectival character of perception take the form:

Perception is perspectival because we perceptually represent certain relational properties in

perspectival similarity as real. The core points in this paper apply even if there is only perspectival variance.

3. Another commonly cited example in spatial visual perception: two same-sized trees at different distances look different and the farther tree looks smaller in some sense—that is, it looks similar to a nearer, smaller tree. See Christopher Peacocke, *Sense And Content: Experience, Thought, And Their Relations* (Oxford: Clarendon Press, 1983).

4. The *lightness* of a surface is the degree to which the surface is white (fully light) or black (fully dark). Lightness is a property of the surface, independent of how it is illuminated.

5. See, for example, David J. Chalmers, “Perception and the Fall from Eden,” in *Perceptual Experience*, ed. Tamar Szabó Gendler and John Hawthorne (Oxford: Clarendon Press, 2006), p. 87.

addition to non-relational properties like size, shape, and color. The relational properties in question are supposed to constitute the varying “looks” or “appearances” that objects can have relative to different conditions of observation. I will reject this sort of explanation and propose that a better explanation takes the form: *Perception is perspectival because perceptual representations have a certain sort of structure.*

I begin in Section 1 by reviewing the arguments that the perspectival character of perception depends on our representation of relational properties. I go on to argue against the two standard approaches to characterizing the relevant relational properties. In Section 2, I discuss the *pluralist approach*. According to the pluralist approach, while the perspectival character of visual spatial perception may depend on our perceptual representations of one kind of relational property, the perspectival character of color perception may depend on our perceptual representations of another kind of relational property, and likewise for hearing and touch. I argue that the pluralist approach is unsatisfying, since it does not account for the seemingly unified nature of the perspectival character of perception. In Section 3, I introduce the *perspectival properties approach*, which holds that all perspectival forms of perception involve representations of a distinctive sort of property: *perspectival properties*. I argue, in Section 4, that this approach violates central commitments of empirical psychology. In Section 5, I offer an alternative type of account. Both the pluralist and perspectival properties approaches attempt to explain the perspectival character of perception in terms of what properties our perceptual states represent. I propose, instead, that the perspectival character of perception depends on the way our perceptual states are structured.

I take for granted that perception is representational. Perceptual states have content about objects in the world and their properties. For example, my current perceptual state represents the tabletop in front of me as brown and rectangular. The content of my perceptual state sets some condition on how the world has to be if that state is to be accurate—my current perceptual state is accurate just in case the tabletop is in fact brown and rectangular. Our perceptual representations exhibit perspectival variance and similarity. A coin has a different look when seen at an angle than when seen head-on in part because our perceptual representations of the slanted coin and of the head-on coin differ in some respect. The slanted coin has a similar look to a head-on ellipse because our perceptual representations of the slanted coin and of the head-on ellipse are similar in some respect.⁶ What features of perceptual representation give rise to variance and similarity?

Let's make this question more precise. We can describe perspectival variance and similarity in terms of the representational features of perceptual states. I will say that a *representational feature* of a perceptual state is a feature that is essential to the way that state represents its subject matter. So my current perceptual state has the representational feature of representing the tabletop. More specifically, this state has the representational feature of representing the tabletop as *brown* and as *rectangular*; and it has this feature because it

6. Some have argued that only the phenomenal character of our perceptual states exhibits variance and similarity. See Peacocke, *Sense And Content*. Those who hold that the representational content of a mental state fully determines that state's phenomenal character (known as *representationalists*) have responded by offering explanations of how perceptual representation is perspectival as well. In this paper, I take for granted that perceptual representation is perspectival and that perspectival variance and similarity arise from how we represent the world. I will leave it open how the perspectival character of perceptual representation relates to the perspectival character of our perceptual phenomenology. I will not comment on the broader question, with which representationalists are concerned, of how the representational content of an experience relates to that experience's phenomenal character.

has the representational feature of being a combination of a representation as of an instance of *brownness* and a representation as of an instance of *rectangularity*. Perspectival variance consists in the fact that one's perception of a slanted coin (or a shaded white surface) and one's perception of a head-on coin (or an unshaded white surface) will have different representational features. Perspectival similarity consists in the fact that one's perception of a slanted coin (or a shaded white surface) and one's perception of a head-on ellipse (or an unshaded gray surface) will share relevant representational features. An explanation of the perspectival character of perception should specify what kinds of representational features account for perspectival variance and similarity.

The standard accounts of the perspectival character of perception suppose that perceptual representations of the head-on coin and the slanted coin have different representational features insofar as they attribute different relational properties to the coin.⁷ Likewise, these views suppose that one's perceptual representations of the slanted coin and the head-on ellipse share representational features insofar as they attribute the same relational properties to the slanted coin and the head-on ellipse. Let's review the arguments for this supposition.

1.1. The Argument from Variance. Consider perspectival variance. A common line of

7. It is not a trivial task to explicate the notion of a relational property. For our purposes, we can consider relational properties as properties whose instantiation in an entity depends metaphysically on how that entity is related to other entities. For example, being slanted and being shaded are relational properties whose instantiation in a surface depends, respectively, on how that surface is related to a viewer and how it is related to an illumination source. I take it that there is a fundamental contrast between non-relational properties such as the orientation and illumination of a surface, and properties such as the shape and color of a surface. Throughout this paper I will describe shape and color as *non-relational* properties—that is, properties whose instantiation in an entity does not depend on how that entity is related to other entities. One may question whether color, or even shape, are truly non-relational. However, little of the substance in this paper hinges on these questions, so long as we have some way of making sense of a basic contrast between properties such as orientation and illumination and properties such as shape and color. For discussions of these issues, see Alex Byrne and David R. Hilbert, “Color Realism and Color Science,” *Behavioral and Brain Sciences* 26, no. 1 (2003): 3–64; Jonathan Cohen, “Color Properties and Color Ascriptions: A Relationalist Manifesto,” *The Philosophical Review* 113, no. 4 (2004): 451–506; Bradford Skow, “Are Shapes Intrinsic?,” *Philosophical Studies* 133, no. 1 (2007): 111–130.

thought goes like this: as you rotate a coin, your perception of the coin changes even as you continue to perceptually represent the coin's circularity—a non-relational property of the coin. Since you continue to perceptually represent the coin as circular in shape, your perception must not be changing with respect to your representation of the coin's shape. But notice that as you rotate the coin, some of its relational properties, such as its orientation with respect to you, change. The change in your perception must, so this line of thought continues, be due to your representing the coin as having different relational properties at different times.

There are some gaps in this argument that need to be filled in. For example, the point is not just that our representation of the coin's circularity remains constant, but that our representation of all of the coin's relevant non-relational properties remain constant. Further, the argument assumes that the representational difference in our perceptions of the head-on and slanted coin is best explained by a difference in what properties are represented. Since you do not represent the head-on and slanted coin as having different non-relational properties, as the argument goes, the difference in your perceptions is that you represent the coin as having different relational properties.

The argument can be formulated more adequately as follows. Let "SLANTED COIN" denote an accurate perceptual representation of the circular surface of a coin that is slanted in depth relative to one's line of sight under normal conditions of observation. Let "HEAD-ON COIN" denote an accurate perceptual representation of the circular surface of the coin were it oriented head-on under normal conditions.

(V1) SLANTED COIN and HEAD-ON COIN have different representational features.

- (V2) The best explanation for why SLANTED COIN and HEAD-ON COIN have different representational features is that they do not represent all the same properties.
- (V3) SLANTED COIN and HEAD-ON COIN represent all the same relevant non-relational properties (for example, circularity).
- (V4) So, the best explanation for why SLANTED COIN and HEAD-ON COIN have different representational features is that there is some relational property that is represented by one but not the other.

Going forward, I will assume V1 and V3 to be true. The difference between one's percepts of the slanted coin and the head-on coin seems to be a difference that bears on how one represents the coin, though it is not a difference in what non-relational properties one represents the coin as having. We can construe V2 as a claim either about what properties are denoted by one's perceptual states or about what properties enter into the perceptual modes of presentation of shape properties.⁸ V2 is perhaps motivated by the idea that the representational features of an experience are entirely determined by what properties the experience represents. According to that view, the only possible explanation for the representational difference between perceptual states is that they represent different properties. But one might find V2 plausible without adopting such a restrictive account of representational features. As a general methodological point, it is often more straightforward to discover and theorize about what objects and properties are psychologically represented than to discover and theorize about other sorts of representational features. Further, one is often, though perhaps not always, experientially aware only of the objects and properties

8. See David J. Chalmers, "The Representational Character of Experience," in *The Future for Philosophy*, ed. Brian Leiter (Oxford University Press, 2004), 153–181; Brad J. Thompson, "The Spatial Content of Experience," *Philosophy and Phenomenological Research* 81, no. 1 (2010): 146–184.

that one perceptually represents. So, it is reasonable to expect that when one is aware of a difference in one's experiences of the head-on coin and the slanted coin, what one is aware of is a difference in the properties possessed by the head-on and slanted coin.⁹ I take V2 to be a reasonable but defeasible methodological assumption.

Notice that perspectival variance does not imply perspectival similarity. In principle your perception of the slanted coin could differ from your perception of the head-on coin without being similar in any respect to our perception of the head-on ellipse. As we will see, some explanations of variance cannot also account for similarity. So we should distinguish between the argument from variance and a parallel argument from similarity.¹⁰

I.2. The Argument from Similarity. The usual formulation of the argument from similarity goes like this: there is some sense in which your perception of the slanted, circular coin is similar to your perception of an appropriately shaped, head-on ellipse—the slanted coin has an “elliptical look,” so to speak. Assuming that you accurately perceive the coin as circular and that your perception is not illusory, you do not also perceive the coin as elliptical. So you do not perceptually attribute the same shape to the slanted coin and head-on ellipse. There must be some relevant relational property that you perceive both the slanted coin and the head-on ellipse as having.

9. See Christopher S. Hill and David J. Bennett, “The Perception of Size and Shape,” *Philosophical Issues* 18 (2008): 294–315.

10. Many philosophers freely shift between talking about variance and similarity. They assume that both perspectival variance and similarity are expressions of the same underlying phenomenon. But not everyone shares this assumption; some acknowledge perspectival variance but deny perspectival similarity. For example, A.D. Smith claims that while one's experience of a coin's shape does differ when one views the coin from different angles, there is no relevant similarity in our perception of a slanted coin and a head-on ellipse (*The Problem Of Perception*, p. 181–2). One cannot use demonstrations of perspectival variance to prove that perspectival similarity is real, as Jonathan Cohen tries to do when, in responding to Smith, he writes: “Like it or not, inconstancy reactions [variance] are part of the observed data; they show that inconstancy [variance] (for example, along the dimensions of size, shape, color) is part of subjects' phenomenal experience in perceiving the cases of interest.” Jonathan Cohen, “Perception and Computation,” *Philosophical Issues* 20 (2010): p. 100.

As with the argument from variance, there are gaps in this reasoning. The argument assumes that representational similarity is best explained by a similarity in what properties one represents. The thought is that because one does not perceptually attribute the same shape, or any other relevant non-relational property, to the slanted coin and the head-on ellipse, it must be that one attributes some relational property to both. Let's fill out the argument. Let "SLANTED COIN" and "HEAD-ON COIN" mean what they did in the argument from variance, and let "HEAD-ON ELLIPSE" denote an accurate perceptual representation of an appropriately shaped elliptical figure, viewed head-on under normal conditions.

- (S1) There is some representational feature that both SLANTED COIN and HEAD-ON ELLIPSE have, but which HEAD-ON COIN does not have.
- (S2) The best explanation for why there is some representational feature that both SLANTED COIN and HEAD-ON ELLIPSE have is that there is some property that they both represent.
- (S3) All the relevant non-relational properties that both SLANTED COIN and HEAD-ON ELLIPSE represent are also represented by HEAD-ON COIN (for example, *having a bounded surface*).
- (S4) So, the best explanation for why there is some representational feature that both SLANTED COIN and HEAD-ON ELLIPSE have, but which HEAD-ON COIN does not have, is that there is some relational property that both SLANTED COIN and HEAD-ON ELLIPSE represent, but which HEAD-ON COIN does not represent.

The arguments from variance and similarity each give us reason to think that an explanation of the perspectival character of perception must involve reference to representations

of relational properties. In the next sections, I discuss different approaches as to what those relational properties are. After considering weaknesses in these approaches, I will recommend that giving up V2 and S2 and pursuing an alternative type of explanation.

II THE PLURALIST APPROACH

Let's begin by considering what I will call the *pluralist approach* to explaining the perspectival character of perception in terms of the representation of relational properties. The pluralist approach applies a divide-and-conquer strategy to explaining why different types of perception are perspectival. A pluralist account might explain the perspectival character of spatial perception in terms of perceptual representations of one kind of relational property, while it might explain the perspectival character of color perception in terms of perceptual representations of another kind of relational property.¹¹ Advocates of the pluralist approach tend to prefer explaining the perspectival character of perception in terms of representational capacities that figure into existing psychological theories and that are mentioned in standard textbooks.¹² I will argue that the pluralist approach's lack of a unified explanation of the perspectival character of perception is unsatisfying.

Consider the slanted coin. What is the representational difference between your perception of the head-on coin and your perception of the slanted coin? One answer that a proponent of the pluralist approach might give is that you represent the head-on coin as

11. Examples of the pluralist approach can be found in Michael Tye, "Perceptual Experience Is a Many-Layered Thing," *Philosophical Issues* 7 (1996): 117–126; Michael Tye, "In Defense of Representationalism: Reply to Commentaries," in *Pain: New Essays on Its Nature and the Methodology of Its Study*, ed. Murat Aydede (The MIT Press, 2005), 163–175; Mohan Matthen, "How Things Look (And What Things Look That Way)," in *Perceiving the World*, ed. Bence Nanay (Oxford University Press, 2010), 226–253; Hopp, "No Such Look."

12. For example, Stephen E. Palmer, *Vision Science* (Cambridge, MA: The MIT Press, 1999); John P. Frisby and James V. Stone, *Seeing: The Computational Approach to Biological Vision*, 2nd ed. (Cambridge, MA: The MIT Press, 2010).

head-on, whereas you represent the slanted coin as *slanted*. It is well established that we perceive how surfaces are oriented with respect to us. So the pluralist approach seems to have a straightforward and plausible explanation for this case of perspectival variance.

While an appeal to the perception of surface orientation can explain perspectival variance in this case, it cannot explain perspectival similarity. Why is your perception of the slanted coin similar to your perception of a head-on ellipse? Under normal conditions, you perceive the slanted coin as slanted and the head-on ellipse as head-on. The similarity in these perceptions cannot be due to your perceptually attributing similar orientations to the coin and the ellipse.

In order to account for perspectival similarity, the proponent of the pluralist approach might appeal instead to representations of the *egocentric directions* to points on a surface—that is, the directions in which the points on a surface are located with respect to one’s viewpoint. It is well-established that one can perceptually represent points in a scene as located at certain distances away from oneself and in certain directions from one’s viewpoint. Now, consider the directions (but not distances) from your point of view to the points on the slanted coin and to the points on a head-on ellipse that perfectly occludes the slanted coin. The directions to the points on the slanted coin will be the same as the directions to the points on the occluding ellipse. To illustrate this, point your finger to the edge of a slanted coin and let your finger trace the outline of the coin. The path of your finger passes through every direction from your shoulder to the outline of the coin. That path—that set of directions—will be the same for an appropriately shaped head-on ellipse. The proponent of the pluralist approach can claim that your perceptions as of the slanted coin and of the head-on ellipse are similar because they both involve representations as of the same set of




Figure 1: The visual angle subtended by two or more points on a surface is the angle α between rays extending from the viewpoint to those points on the surface. The visual angle subtended by two points on a surface depends on the distance S between those and the distance D between the object and the viewpoint.

directions from your viewpoint to the points on the represented surface.

In fact, this proposal needs a slight modification. Suppose the slanted coin is adjacent to the head-on ellipse. The set of absolute directions from you to the points on the two surfaces will be different because one set of directions will point toward the coin and the other toward the ellipse. The explanation should appeal instead to what is in common between these sets of directions—namely, the *relations* among the directions. The *visual angle* of a surface corresponds to the relevant relations among the directions from one’s viewpoint to points on the surface of an object. Roughly, the visual angle of a pair of visible points on an object is the angle or difference between the directions from a point of view to those two points (Figure 1). The visual angle of a surface as a whole corresponds to the set of visual angles between every pair of visible points on the surface, or to the shape of a cone whose base is the visible portion of the surface and whose apex is one’s viewpoint.

The head-on coin has a different visual angle than the slanted coin, which in turn has the same visual angle as a head-on ellipse. So, representations of the visual angles of surfaces are well-suited to account for both perspectival variance and similarity in size and shape perception.¹³ Moreover, there is some precedence in the vision science literature for the

13. See Tye, “Perceptual Experience Is a Many-Layered Thing”; René Jagnow, “Representationalism and the Perspectival Character of Perceptual Experience,” *Philosophical Studies* 157, no. 2 (2012): 227–249.

claim that we perceptually represent visual angles.¹⁴ A representation of the visual angle between two (or more) points in a scene could subserve the perception of size and even guide actions by specifying, for example, how much one's eye or head would have to rotate in order to shift focus from one point to the other.

Representations of visual angles cannot, however, explain other cases of perspectival variance and similarity. For example, one cannot appeal to the perception of visual angles in order to explain why there is some respect in which one's perception of the shaded side of a uniformly painted wall is different than one's perception of the unshaded side or why there is some respect in which it is similar to one's perception of an unshaded surface with a darker paint. The proponent of the pluralist approach will have to find some other way of explaining these patterns in color and lightness perception—and likewise for cases in hearing, touch, and so on. A proponent of the pluralist approach may, for example, have to posit representations of surface luminance in order to explain perspectival variance and similarity in color and lightness perception.¹⁵

The pluralist cannot give a unified explanation of the perspectival character of perception. According to the pluralist approach, perception is perspectival because of what relational properties we perceptually represent; but the relevant properties are fundamentally different for different forms of perception. Visual angle and surface luminance, for instance, do not form a natural kind. So it cannot be in virtue of representing a common kind of property that these representations form a unified psychological kind. Substantially

14. See Don McCready, "On Size, Distance, and Visual Angle Perception," *Perception & Psychophysics* 37, no. 4 (July 1985): 323–334; Hirohiko Kaneko and Keiji Uchikawa, "Perceived Angular and Linear Size: The Role of Binocular Disparity and Visual Surround," *Perception* 26, no. 1 (1997): 17–27.

15. A surface's *luminance* is the total light energy reflected by that surface—a product of the surface's reflectance and the way it is illuminated.

different sensory cues and computations would be responsible for forming representations of visual angles than would be responsible for forming representations of surface luminance, and these representations would have substantially different kinds of influences on other psychological capacities. The pluralist cannot say what underlying nature these representations share.

The disunified character of the pluralist approach is unsatisfying. It is striking that while there are many differences between spatial perception and color perception, they both exhibit patterns of variance and similarity. Perspectival representations seem to constitute a robust psychological kind. The pluralist approach purports to explain different cases of perspectival variance and similarity, but it does not offer a systematic, common account of these cases. All things being equal, we should look for a more unified, systematic account.¹⁶

III THE PERSPECTIVAL PROPERTIES APPROACH

In contrast to the pluralist approach, what I call the *perspectival properties approach* offers a unified account of the perspectival character of perception. This approach posits the existence of a set of distinctive *perspectival properties* and holds that all perception involves the representation of such properties. Perspectival properties are supposed to constitute a unified natural kind of relational property and, correspondingly, representations of these

16. I have emphasized the disunified nature of the pluralist approach. There are other worries about attempts to explain the perspectival character of shape and color perception in terms of representations of visual angle and surface luminance, respectively. In the first place, it may be that while visual angle and surface luminance are sensorily registered, they are not represented in perception. See Tyler Burge, *Origins of Objectivity* (New York: Oxford University Press, 2010). A separate worry is that representations of visual angle and surface luminance, insofar as they are accurate, would not account for the patterns of perspectival variance and similarity that are found in perception. For example, in “The Perception of Size and Shape,” Hill and Bennett argue that perspectivally similar shape representations are often not of surfaces that have similar visual angles. While these concerns are important, I have set them to the side in order to focus on broader explanatory issues.

properties are supposed to constitute a unified psychological kind.¹⁷ Unlike the sorts of representations to which the pluralist appeals, representations of perspectival properties do not resemble the kinds of representations that ordinarily figure into empirical models of perception. In the next section, I will argue that the perspectival properties approach is not a tenable empirical hypothesis and that there is good reason that perceptual psychology standardly does not posit representations of perspectival properties.

Philosophers have characterized perspectival properties in different ways. Perspectival properties have been characterized, for example, in terms of the optical projections that objects have or else in terms of the perceptual states that objects are disposed to cause.¹⁸ Different versions of the perspectival properties approach hold that perceptual representations of non-perspectival properties depend either epistemically, computationally, or constitutively on perceptual representations of perspectival properties.¹⁹

In order to evaluate formulations of the perspectival properties approach empirically, one must have some conception of the conditions under which something instantiates one perspectival property or another. One needs at least a tentative account of the identity conditions of determinate perspectival properties in order to ascertain whether a representation of a perspectival property is accurate. And knowing whether a perceptual representation of

17. Examples of the perspectival properties approach can be found in Michael Huemer, *Skepticism and the Veil of Perception* (New York: Rowman & Littlefield, 2001); Alva Noë, *Action In Perception* (Cambridge, MA: MIT Press, 2004); Susanna Schellenberg, “The Situation-Dependency of Perception,” *The Journal of Philosophy* 105, no. 2 (2008): 55–84; Cohen, “Perception and Computation”; Berit Brogaard, “Strong Representationalism and Centered Content,” *Philosophical Studies* 151, no. 3 (2010): 373–392; Christopher S. Hill, “The Content of Visual Experience,” in *Meaning, Mind, and Knowledge* (New York: Oxford University Press, 2014), 218–236. Tye is sometimes grouped together with proponents of the perspectival properties approach. However, since Tye explains perspectival phenomena in visual spatial perception in terms of representations of visual angles while explaining perspectival phenomena in color perception in terms of representations of illumination, I think it is better to classify him as a pluralist.

18. See, respectively, Noë, *Action In Perception*; Cohen, “Perception and Computation.”

19. See, respectively, Schellenberg, “The Situation-Dependency of Perception”; Cohen, “Perception and Computation”; Noë, *Action In Perception*.

a perspectival property is accurate or not is critical for developing and testing computational theories of how representations of perspectival properties are generated and employed. The claim that we represent *perspectival ellipticality*, for example, and the corresponding question of how one's visual system computes a representation of *perspectival ellipticality* are not well defined until there is at least a preliminary conception of what *perspectival ellipticality* is.

For purposes of illustration, I will describe a *projective* model of perspectival properties. This model will help to illustrate the advantages and disadvantages of the approach. I will say that an object instantiates a particular *projective property* by virtue of projecting a certain pattern onto a projection plane relative to a viewpoint. To be *projectively elliptical* relative to a projection plane and viewpoint is to project an elliptical pattern onto the projection plane relative to that viewpoint. I will assume that the perspectival properties approach only claims that we represent items as having projective properties that correspond to the projected properties registered by the perceptual system. In the case of vision, the relevant definitions of "projection," "projection plane," and "viewpoint" will be specified by reference to laws of optics and the anatomy of the eye. Let's think of the "projection plane" and "viewpoint" as (possibly empty) regions of physical space.²⁰ So an object's projected image can be specified as a cross-section, at a projection plane, of the light rays that the object reflects and that converge at the specified viewpoint. In canonical form:

Projective Property An object x is projectively- F relative to a viewpoint v and a projection plane p iff the optical projection of x onto p relative to v has the property of

20. For concreteness, I am characterizing projective properties in terms of optical projections at the surface of the retina. Projective properties could alternatively be specified, for example, with respect to the "Cyclopean Eye" located midway between the two actual eyes. The choice of the viewpoint and projection plane will not matter to the current discussion.




Figure 2: The circular disk in (a) and the ellipse in (b) are both projectively elliptical relative to their respective viewpoints and projection planes.

being F .

As an example, consider the circular disk in Figure 2a. This disk is projectively elliptical with respect to the illustrated projection plane and viewpoint. The ellipse in Figure 2b is also projectively elliptical with respect to a projection plane and viewpoint.²¹

I will assume that objects have projective properties. One should be careful not to confuse the *project-ive* property with the property of the projected image (the *project-ed* property). The circular surface of a disk is *projectively elliptical* relative to the appropriate viewpoint; the pattern that the disk’s surface projects to the projection plane is simply *elliptical*. Projective shape, as I have defined it, is a property of the distal object, not of that object’s projection. One should also be careful to distinguish between the projective properties that a thing instantiates and the projective properties that a perceiver represents. While an object will simultaneously instantiate an enormous number of projective shapes, the perspectival properties approach only claims that we perceptually represent things as having the projective properties that correspond to the projected stimuli that our sensory

21. It will be convenient to suppress reference to viewpoints and projection planes and just say that an object is “projectively- F ,” full stop, when there is some salient viewpoint v and projection plane p (as determined, for example, by the perceiver’s eye or visual system) such that the projection of x onto p relative to v has the property F . Two objects can share projective properties, in this sense, relative to different viewpoints and projection planes. For example, the disk and the ellipse in Figure 2 are both *projectively elliptical*, full stop.

systems are registering at a given time.

The projective characterization of perspectival properties must be treated cautiously if it is to be presented in the best light. According to the present account, the coin is perspectivally elliptical if and only if its projection on a plane relative to a viewpoint is elliptical. It might be tempting to express the fact that a subject is representing the coin as *perspectivally elliptical* by saying that the subject represents it as *having an elliptical optical projection on a plane relative to a viewpoint*. But in order to represent the coin as *having an elliptical optical projection on a plane relative to a viewpoint*, one would seemingly have to be able to perceptually represent points in space as *viewpoints*, regions in space as *projection planes*, and one would have to perceptually represent the rules of optical projection. While one may perceptually represent a point in space (for example, a point midway between the eyes) as one's viewpoint, there is little evidence to support the claim that one can perceptually represent as of projection planes or the rules of optical projection. The perspectival properties approach loses some plausibility if it is committed to one's representing these things. But it need not have this commitment.

To see how the perspectival properties approach can avoid such commitments, we should distinguish between representing a perspectival property and representing its identity conditions. The identity conditions for a perspectival property are specified by our theory. We do not need to represent a property's identity conditions in order to represent that property. So, while the perspectival properties approach holds that we represent the slanted coin as *projectively elliptical*, it need not be committed to the implausible view that we represent the coin as *having an elliptical optical projection on a plane relative to a viewpoint*.²²

22. The same issues arise with non-projective characterizations of perspectival properties. For example,

The unity of the perspectival properties approach can be brought out by looking at the difference in explanatory power between projective properties and visual angles. These two sorts of properties are often treated interchangeably.²³ But the concept of a projective property has a generality that the concept of visual angle does not. For any property of a projection, we can define a corresponding projective property. We introduced the concept of projective ellipticality, for example, by reference to our ordinary concept of ellipticality, applied to the image on a projection plane. By turning to other features of the projected image, we can extend the notion of a projective property from the geometrical cases we have been focusing on, such as projective shape, to non-geometrical cases such as projective color. Projective color can be specified in terms of color, or some related property, at a projection plane. With a general enough concept of projection we can even define non-visual projective properties, such as the projective loudness of a sound. By contrast, the notion of a visual angle cannot be extended in any natural way to account for perspectival color or loudness perception. Because we can plausibly identify projective properties for all forms of perception, and since these properties all have the same canonical nature, one could offer a unified theory of the perspectival character of perception by appealing to representations of projective properties.

Cohen gives a dispositional, rather than projective, account of perspectival properties, and calls them “perceptual state dispositions.” He writes that when we look at the slanted coin, the visual system “represents the distal item as bearing this perceptual state disposition: *disposed to generate in us an instance of the type of perceptual state we undergo when perceiving an ellipse straight on*” (“Perception and Computation,” p. 110). Cohen’s formulation suggests that the perceptual system represents as of *perceptual states, dispositions, causal generation*, and so on. But there is no independent empirical reason to think that the perceptual system can represent as of such properties and relations. See Tyler Burge, “Vision and Intentional Content,” chap. 14 in *John Searle and His Critics*, ed. Ernest Lepore and Robert van Gulick (Oxford: Blackwell, 1991), 195–213.

23. See, for example, Huemer, *Skepticism and the Veil of Perception*; David J. Bennett, “Varieties of visual perspectives,” *Philosophical Psychology* 22, no. 3 (2009): 329–352.

We now have two candidate approaches to explaining the perspectival character of perception. The pluralist approach is unsatisfying because it does not offer a unified account of the perspectival character of perception. The perspectival properties approach, by contrast, offers a unified alternative. I will now argue that the perspectival properties approach violates a central commitment of empirical psychology, since the perceptual representations that it posits are functionally redundant. For concreteness, I will focus on a projective characterization of perspectival properties.

Moving forward, I will assume that our perceptual capacities depend on information-processing operations that generate representations of how the world is on the basis of sensory stimuli. This assumption is at the heart of what is now the standard approach to studying perception in psychology.²⁴ A central commitment of this *information-processing approach* in psychology is that all psychological states make some contribution or other to information processing. For example, psychological states may carry new information or content, contribute to the production of new information or content, or help maintain information or content that is already available. From the perspective of the information-processing approach, we should not posit functionally redundant psychological states—states that play no role in information processing. This does not mean that we should never posit psychological states that play less than optimal roles in information processing. Nor does it mean that we should never posit states whose information-processing roles are not in

24. See David Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (Cambridge, MA: The MIT Press, 1982); Stephen E. Palmer and Ruth Kimchi, “The Information Processing Approach to Cognition,” in *Approaches to Cognition: Contrasts and Controversies*, ed. T.J. Knapp and L.C. Robertson (Hillsdale, NJ: Lawrence Erlbaum Associates, 1986), 37–77.

general conducive to the survival and reproduction of the perceiver. Finally, it does not mean that we should never posit multiple psychological states that have the same or similar contents. Rather, the claim is that we should not posit psychological states that would not make any identifiable contribution whatsoever in psychological information processing. A fundamental assumption of contemporary psychology is that perceptual states must have something to do—whether or not they do it well and whether or not their doing it is good for the organism’s well-being.

The perspectival properties approach claims that we perceptually represent perspectival properties. For this to be a tenable hypotheses, from the perspective of the information-processing paradigm in psychology, representations of perspectival properties would have to make some identifiable contribution to perceptual information processing. In fact, some proponents of the perspectival properties approach hold that representations of perspectival properties have an important role to play in perceptual processing.²⁵ They maintain that the perceptual system computes representations of non-perspectival properties partly on the basis of prior representations of perspectival properties. On this account, when we view the slanted coin the visual system first registers an elliptical pattern on the retina, produces a representation as of a perspectivally elliptical surface, and then on the basis of this transitional representation produces representation as of a circular surface. The idea is that representations of perspectival properties are stepping stones on the way to representing non-perspectival properties.

However, the claim that representations of perspectival properties are stepping stones to representing properties like size, shape, and color is untenable. I will argue that the

25. See, for example, Cohen, “Perception and Computation.”

perceptual system is set up so that it can go from registering the sensory stimulus to representing non-perspectival properties such as shape without any intermediate representations of non-perspectival properties. The regularity that objects instantiate the relevant perspectival properties is reflected in the architecture of the perceptual system, and so it does no good to also represent particular objects as instantiating those properties on particular occasions. More generally, there is no identifiable contribution that representations of perspectival properties would make to perceptual processing.

My argument rests on a standard understanding within perceptual psychology of how the perceptual system takes advantage of environmental regularities.²⁶ The task of perception is to represent the distal stimuli that gave rise to the proximal impacts on our sensory receptors. In principle, infinitely many distal features could have given rise to a given proximal stimulus. The task of perception is tractable, however, so long as only a subset of the distal features that *could* in principle give rise to a particular proximal stimulus *do* give rise to it and if there are regularities constraining which distal features give rise to which proximal stimuli. Evidently, the perceptual system takes advantage of this fact. The perceptual system operates as if certain kinds of regularities—called “natural constraints”—govern the way distal features give rise to proximal stimuli. Among these natural constraints are the regularities that light travels in straight lines; that light comes from above; that most points lie on rigid surfaces; for some animals, that the period of the Earth’s rotation is about twenty-four hours; and so

26. See Shimon Ullman, “The Interpretation of Structure from Motion,” *Proceedings of the Royal Society of London. Series B. Biological Sciences* 203, no. 1153 (1979): 405–426; Marr, *Vision*; Donald D Hoffman and Whitman A Richards, “Parts of Recognition,” *Cognition* 18, no. 1 (1984): 65–96; Roger N. Shepard, “Perceptual-Cognitive Universals as Reflections of The World,” *Behavioral and Brain Sciences* 24, no. 4 (2001): 581–601; Laurence T. Maloney, “Surface Colour Perception and Environmental Constraints,” chap. 9 in *Colour Perception: Mind and The Physical World*, ed. Rainer Mausfeld and Dieter Heyer (New York: Oxford University Press, 2003), 279–300; Zenon Pylyshyn, *Seeing and Visualizing: It’s Not What You Think* (Cambridge, MA: The MIT Press, 2003).

on. Perception fails to be veridical when such regularities do not obtain.

Perceptual systems operate as if some (but perhaps not all) of these environmental regularities or natural constraints are constants that hold always and everywhere. For example, the visual system operates as if all light travels in straight lines. The standard view of information-processing theories is that the perceptual system treats these regularities as constants in virtue of the way its *architecture* is set up, where the architecture of the perceptual system consists in the information-processing operations that it carries out. Through either evolution or individual development within an environment that satisfies certain regularities, the procedures carried out by one's perceptual system produce, as a matter of course, representations that would be approximately veridical were those regularities to obtain. I will paraphrase this situation by saying that these natural constraints are “reflected” in the processing architecture of the perceptual system—in the way the information-processing operations in the perceptual system make certain perceptual outputs possible (or likely), and others impossible (or unlikely), for a given sensory input.²⁷ To be clear, when I say that the architecture “reflects” a natural constraint I do not mean that the constraint is specified in the content of some representation in the system. Rather, the system is just set up to operate as if the constraint holds.²⁸

The idea that certain regularities may be “reflected” in the processing architecture of a system may be clearer with a rough analogy. Suppose one wants to build a *times-seven* device that will multiply a given number by seven. Since the multiplier, seven, is always the

27. It is worth noting that whether a natural constraint is reflected in the architecture of the perceptual system is independent of whether the relevant aspects of the system's architecture are innate or learned, fixed or changeable.

28. See Michael Kubovy and William Epstein, “Internalization: A metaphor we can live without,” *Behavioral and Brain Sciences* 28, no. 4 (2001): 618–625; Pylyshyn, *Seeing and Visualizing: It's Not What You Think*; Burge, *Origins of Objectivity*.

same, the simplest way to build the device is to “hardwire” or “gear” it so that for any input representation of a number it automatically returns a representation of that number times seven. For example, suppose we represent positive numbers by the number of rotations that a gear makes. Our device may consist of two gears having the appropriate ratio of teeth so that every full rotation of the larger input gear translates into seven full rotations of the smaller output gear. The multiplier, seven, is reflected in the relative sizes of the gears. It is important to note that the device’s ability to multiply numbers by seven does not depend on its occasioning a representation of the multiplier in the turns of a particular gear.

Here is the first central point of my argument against the claim that we represent perspectival properties. The architecture of the visual system reflects the natural constraint that objects instantiate projective properties corresponding to the patterns registered on the retina. In fact, some of the most basic natural constraints on visual perception come from the laws of optics, or some approximation of them. It is a widespread commitment of empirical models in perceptual psychology that the architecture of our visual system is set up to operate as if something like the laws of optics always govern the formation of the inputs to vision.²⁹ So, given an elliptical retinal image, for example, a normal visual system could only output representations of shapes (circles and ellipses) that could have projected that input at some orientation or other. A normally functioning visual system could not return a representation as of a square in response to an elliptical pattern on the retina.

Now for the second central point of the argument. There is no functional benefit to

29. Indeed, perceptual psychologists often are interested in more substantial natural constraints that either entail or rely on the laws of optics—for example, the constraint that collinear lines in the retinal image are only projected by collinear contours in the scene, or that points of deep concavity in the retinal image are only projected by points at which one object either intersects or occludes another. Psychologists typically suppose that the architecture of the visual system reflects even these natural constraints, which are in fact less universal than the laws of optics. See, for example, Hoffman and Richards, “Parts of Recognition.”

representing what is already reflected in a system's architecture. The visual system does not, for example, need to represent the projective ellipticality of the slanted coin in order to represent that coin's circularity, since the visual system is already constrained by its architecture to only generate representations of those shapes that could have projected the elliptical pattern at the eye. The system can go from registering the proximal stimulus to representing the non-projective properties of the distal stimulus, without any intermediate representations of projective properties.

Consider, by analogy, our times-seven device. Since the multiplier, seven, is reflected in the relative sizes of the input and output gears, there is no functional benefit to including a separate gear whose rotations represent the multiplier. The representation of the multiplier is not needed as an intermediate step in generating the right output. Likewise, if the processing architecture of the perceptual system reflects the constraint that perceived objects project the registered proximal stimuli, then there is no functional benefit to also tokening states that represent those particular objects as projecting the registered stimuli. Tokening such states would not contribute anything over and above what is already guaranteed by the architecture of the perceptual system.

Since representations of projective properties would be functionally redundant, we should not posit them. The dominant information-processing paradigm in psychology assumes that psychological states must make some sort of contribution to information processing. But since the architecture of the perceptual system already reflects the constraint that the objects of perception project the registered sensory inputs, perceptual representations of those objects as instantiating projective properties would not contribute anything to perceptual information processing. The perspectival properties approach violates a foundational com-

mitment of perceptual psychology by positing representations that would make no identifiable contribution to perceptual information processing.

Non-projective variants of the perspectival properties approach face versions of this same argument. Cohen, for example, offers a dispositional account on which to be *perspectivally elliptical* is to be *disposed to generate in us an instance of the type of perceptual state we undergo when perceiving an ellipse straight on*. But the architecture of the perceptual system already reflects the constraint that the objects of perception are disposed to be the distal causes of the system's states. It is natural that the perceptual should reflect this constraint because, plausibly, the primary and unchanging function of the visual system is to represent the distal causes of its states. If the perceptual system reflects this constraint, then there is no reason for the system to occasion representations of particular objects as *disposed to cause such-and-such states of the perceptual system*.³⁰

The perspectival properties approach promises a unified account of the perspectival character of perception, but at the expense of positing perceptual representations that would serve no role in perceptual information processing. If a unified account is achieved only by appealing to states that have no apparent function in our psychologies, then the unification is not worth it.

V THE STRUCTURES OF PERSPECTIVAL REPRESENTATIONS

The view that the perspectival character of perception depends on our perceptually representing certain relational properties standardly takes the form of either a disunified, pluralist account or an account that posits functionally redundant psychological states. I will now

30. See Burge, "Vision and Intentional Content."

recommend a different sort of approach that has neither of these drawbacks. According to what I call the *structural approach*, the central difference between perspectively variant representations is that they are structured in different ways from their parts, while the central commonality between perspectively similar representations is that they are structured in similar ways from their parts. After introducing the notion that perceptual states have part-whole structure, I will illustrate how differences and similarities in this kind of structure can account for perspectival variance and similarity. I will then argue that what unifies the different systems of perspectival representation is the way in which, in those systems, representations of properties such as size, shape, and color are structurally interlocked with representations of properties such as distance, orientation, and illumination.

A basic assumption in perceptual psychology is that representational perceptual states—the vehicles of perceptual content—have part-whole structure. Perceptual states can, so to speak, be “made of” or arranged from other constituent perceptual states, much as sentences are made of component words and maps are made of colored regions representing territories. For example, my representation of the tabletop as brown and rectangular has as constituent parts my representation as of an instance of brownness and my representation as of an instance of rectangularity. The structure of a representation consists in the way its constituent parts are combined.³¹ The information-processing architecture of a system is intimately related to the structure of the representations over which the system operates.

Some of the core processes in perception involve building up, breaking down, searching, or

31. Throughout, I will be speaking of the part-whole structure of the vehicles of representational content. But it may be helpful to think of a representation’s content as itself having a part-whole structure that mirrors the structure of the vehicle, by analogy with the idea that sentences have *structured meanings* that reflect the syntax of those sentences. See David Lewis, “General Semantics,” *Synthese* 22, nos. 1/2 (1970): 18–67; M.J. Cresswell, *Structured Meanings* (Cambridge, MA: The MIT Press, 1985); Scott Soames, “Direct Reference, Propositional Attitudes, and Semantic Content,” *Philosophical Topics* 15, no. 1 (1987): 47–87.

comparing the structures of representations.³²

The structure of a perceptual state is an important representational feature of that state. In general, the structure of a representation plays an essential role in its representing what it does. What something represents normally depends on what its constituents represent and how they are combined. What state of affairs a sentence represents, for example, depends on the individuals and properties that the words in that sentence denote and how those words are put together. What a map depicts depends on what the colors in that map depict and how the colored regions of the map are arranged. On the other hand, a representation's structure is not fully determined by what it represents. A map and a (very long) sentence could represent all the same locations and relations among all the same landmarks. Yet these representations have different structures, being made up in different ways from different sets of primitive representations.

How might the structure of a representation account for its perspectival character? Consider paintings made in a realistic style. Such paintings are paradigmatic cases of perspectival representations. Intuitively, such paintings exhibit perspectival variance and similarity because of how they are structured—how they are composed from colored marks on a two-dimensional surface. The pattern you paint in order to depict a slanted coin has to be different from the pattern you use to depict a head-on coin (perspectival variance) and will

32. For classic discussions of how psychological representations are structured and how their structure relates to information-processing architectures, see Stephen E. Palmer, "Fundamental Aspects of Cognitive Representation," in *Cognition and Categorization*, ed. Eleanor Roach and Barbara B. Lloyd (Hillsdale, NJ: Lawrence Erlbaum Associates, 1978), 259–303; Stephen Michael Kosslyn, *Image and Mind* (Cambridge, MA: Harvard University Press, 1980); Marr, *Vision*; Zenon W. Pylyshyn, *Computation and Cognition: Toward a Foundation for Cognitive Science* (Cambridge, MA: The MIT Press, 1984).

Psychologists sometimes use the term "format" to refer generally to the available ways of structuring representations in a system. Sentences and maps have different formats, in this sense, while a map of South America and a map of Africa may have the same format but different structures.




Figure 3: Rough illustrations of array representations as of (a) a head-on circular surface, (b) a slanted circular surface, and (c) a head-on elliptical surface. The primitive constituents are numerals that represent contour segments as being located in a particular distance from the viewpoint. The direction in which we represent a contour segment as located is a function of where the numeral is placed in the array. Dotted lines indicate which sets of primitives integrate into complex representations. These complex representations attribute features like size and shape to whole surfaces, depending on what distances their primitive parts represent and how those parts are placed in the array.

be similar to the pattern you paint to depict a head-on ellipse (perspectival similarity). The paint you use to depict a shaded white surface has to be different from the paint you use to depict an unshaded white surface (perspectival variance) and similar to the paint you use to depict an unshaded gray surface (perspectival similarity). There is no pressing need to hypothesize that paintings depict visual angles, surface luminance, or perspectival properties. The perspectival character of paintings is a product of their structure.

We can turn to some empirical hypotheses about the structure of perceptual representations to illustrate how the structures of perceptual states accounts for the perspectival character of perception—though the structural approach need not be tied to these specific hypotheses. On many accounts, perceptual representations of surfaces are structured like arrays, analogous to those in Figure 3.³³ To say that a perceptual representation is array-like

33. See, for example, Marr, *Vision*; Gareth Evans, “Molyneux’s Question,” in *Collected Papers* (Oxford: Clarendon Press, 1985), 364–399; Christopher Peacocke, *A Study of Concepts* (The MIT Press, 1992); Tyler Burge, “Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints,” *Philosophy and Phenomenological Research* 88, no. 2 (2014): 485–501.

The literature on *mental imagery* contains many discussions about the concept of an array-like representation. See Kosslyn, *Image and Mind*; Steven Pinker, “A computational theory of the mental imagery medium,” in *Cognitive and neuropsychological approaches to mental imagery*, ed. Michel Denis, Johannes Engelkamp, and John T.E. Richardson (Martinus Nijhoff Publishers, Boston), 17–32; Michael Tye, *The*

is not to say it is literally laid out in space, nor does it suggest that we have some “inner eye” that views and interprets the patterns in an “inner array.”³⁴ Rather, the view is that we can best explain how certain perceptual representations are processed and how they can combine with each other in terms of what cells in an array the constituents of those representations occupy. The constituents of array-like perceptual representations are perceptual states that represent patches and edges of surfaces as having certain orientations and distances from us. The cell or address that a constituent representation has in the representational array, like a pixel’s place in a digital photograph, corresponds to the line of sight along which that constituent represents a patch of surface—with adjacent cells corresponding to adjacent lines of sight. An array-like representation of a surface is a combination of representations in these cells. It is plausible that array-like perceptual representations underlie important aspects of conscious visual experience.³⁵

If perceptual representations of surfaces are array-like, then the representation of the head-on coin’s circularity must have a different structure than the representation of the slanted coin’s circularity. The representations of the head-on coin and the slanted coin must be organized over different cells in the representational array, since the head-on coin occupies different lines of sight than the slanted coin (compare with Figure 3a and b). This structural difference accounts for perspectival variance in our perception of the coin.

On the other hand, the representation of the slanted coin is structurally similar to the

Imagery Debate (The MIT Press, 1991).

34. See Ned Block, “Mental Pictures and Cognitive Science,” *The Philosophical Review* 92, no. 4 (1983): 499–541.

35. For arguments to this effect, see Ray Jackendoff, *Consciousness and the Computational Mind* (Cambridge, MA: The MIT Press, 1987); Jesse J. Prinz, “The Intermediate Level Theory of Consciousness,” in *The Blackwell Companion to Consciousness*, ed. Max Velmans and Susan Schneider (Malden, MA: Blackwell Publishing, 2007), 248–260; David J. Bennett, “Seeing Shape: Shape Appearances and Shape Constancy,” *The British Journal for the Philosophy of Science* 63, no. 3 (2012): 487–518.

representation of the head-on ellipse. To be sure, the two representations will have different constituents. The representation of the slanted coin will be organized from representations as of slanted surface patches, whereas the representation of the head-on ellipse will be organized from representations as of head-on surface patches. It is because the perceptual states that represent the slanted coin and the head-on ellipse have different sets of constituents that the one state represents the coin as slanted and circular and the other state represents the ellipse as head-on and elliptical. However, while the two perceptual states represent different shapes and are arranged from different set of constituents, those sets of constituents are organized over the same pattern of cells in the representational array, since we see the points on the slanted coin and the head-on ellipse along all the same lines of sight (compare with Figure 3b and c). This structural similarity accounts for the perspectival similarity of our perceptions of the slanted coin and the head-on ellipse.³⁶

While the structures of the representations involved in lightness and color perception are in some ways more difficult to describe, there is good reason to think that they are responsible for the perspectival phenomena we find there. Plausibly, how one represents a surface's lightness and color depends on how one represents that surface as illuminated.³⁷

36. In "Perceptual Experience Is a Many-Layered Thing," Tye draws a similar connection between the array-like structure of our perceptual representations and the impression that a farther tree looks different and, in some sense "smaller," than a nearer tree of the same size. Tye suggests that array-like perceptual states represent visual angles. On Tye's final analysis, our perception of the farther tree is perspectivally different from our perception of the nearer, same-sized tree because we attribute different visual angles to the trees; the distant tree is perspectivally similar to a nearer, smaller tree because we attribute the same visual angles to these trees. In contrast to Tye, I do not claim that array-like perceptual states represent visual angles. I believe that even if a representation's place in the representational array corresponds to the line of sight along which a point in the scene is represented, combinations of these representations need not represent as of the visual angles between those lines of sight. More importantly, however, I claim that a unified account of perspectival variance and similarity should center on the *structure* of our perceptual states and not on what properties, such as visual angle, those states might represent.

37. While I think it is plausible that we represent how surfaces are illuminated there is no current consensus on the matter. Many models of lightness and color perception assume that the visual system merely filters out the effect of illumination without representing that effect. If we do not represent something like how surfaces are illuminated, then the structural approach will have to be pursued along different lines than I suggest

Light at the retina is the product of both the material properties of the surfaces in front of us and of the way those surfaces are illuminated. A task of color vision is to decompose the light at the retina into material lightness and color components and an illumination component. The visual system seems to reflect the regularity that the represented material components and the represented illumination components jointly account for the light registered at the eye. This means that represented lightness and color and represented illumination have a complementary relationship. If, given the light at the eye, you can represent a surface as a well-lit gray or else as a shaded white, you could not normally represent that surface on that occasion as a shaded gray or a well-lit white. This regularity may well be embodied in the very structure of representations of lightness and color. One can then think of the registered light at the eye as setting a structural constraint on how representations of lightness and color can combine with representations of illumination. For example, a representation of a surface as shaded and white or a representation of the surface as unshaded and gray may both be possible under a particular structural constraint, while a representation of the surface as shaded and gray would not be possible under that constraint. If this is right, then one's representations of the shaded white surface and of the unshaded white surface will differ because they are subject to different structural constraints. And one's representations of the shaded white surface and of the unshaded gray surface will be similar because they fall

here. For some discussion of these issues, see Laurence T. Maloney and Joong Nam Yang, "The Illumination Estimation Hypothesis and Surface Color Perception," chap. 11 in *Colour Perception: Mind and The Physical World*, ed. Rainer Mausfeld and Dieter Heyer (New York: Oxford University Press, 2003), 335–358; Alan Gilchrist, *Seeing Black and White* (New York: Oxford University Press, 2006); Frederick A.A. Kingdom, "Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy," *Vision Research* 51, no. 7 (2011): 652–673.

under a similar structural constraint.³⁸

I have been proposing that the perspectival character of size, shape, lightness, and color perception depends on the structures of the representations involved. I now want to suggest that all types of perspectival variance and similarity result from a common kind of structural characteristic: in all perspectival forms of perception, representations as of non-relational properties (for example, size, shape, surface lightness, color) and representations of relational properties that characterize how things relate to us and to their surroundings (for example, distance, direction, orientation, illumination) are *structurally interdependent*. I will say that one type of representation α (for example, a shape representation) structurally depends on another type of representation β (for example, an orientation representation) if instances of β constrain the instances of α with which they can combine.

Paintings offer good examples of this structural interdependency. The way paintings are put together by combining colored marks on a two-dimensional surface requires that the types of marks you use to depict shape and lightness/color intrinsically depend on the types of marks you use to depict orientation and illumination, respectively. In principle, there are only certain ways you can depict the shape of a surface given that you are depicting the surface at a certain orientation, and *vice versa*. This structural interdependency gives rise to perspectival variance and similarity in paintings. Because the depiction of orientation structurally constrains the depiction of shape, you cannot paint a slanted disk the same way that you paint a head-on disk and you must paint a slanted disk the same way you paint a head-on ellipse. Likewise, how you depict a surface's color depends on how you depict the

38. See also Rainer Mausfeld, "Conjoint Representations and the Mental Capacity for Multiple Simultaneous Perspectives," in *Looking Into Pictures: An Interdisciplinary Approach to Pictorial Space*, ed. Heiko Hecht, Robert Schwartz, and Margaret Atherton (Cambridge, MA: The MIT Press, 2003), 17–60.

way it is illuminated, and *vice versa*. As a result of this interdependency, you cannot paint a shaded white surface the way you paint an unshaded white surface and the way you paint a shaded white surface must be similar to the way you paint an unshaded gray surface.

The sort of structural interdependencies in painting show up throughout perception. For example, what structure an array-like representation of a circular surface can have depends on how we represent the surface as oriented. Conversely, how we represent a surface as oriented depends in part on the structure of our array-like representation of that surface. In the case of color perception, it is plausible that one's representation of the way a surface is illuminated constrains what color representations are structurally possible, and *vice versa*. Perspectival variance and similarity arise out of the specific ways in which representations of non-relational properties such as size, shape, lightness, and color and representations of relational properties such as distance, orientation, and illumination, are interdependent.

In my objection to the perspectival properties view, I argued that there is no need to represent the objects of perception as producing the registered proximal stimuli because that they do so is reflected in the architecture of the perceptual system. This point is not incidental to the current proposal. If certain constraints are reflected in the very architecture of the system, then we have reason to expect those constraints to be embodied in the structures of the representations in which the system traffics. For example, if the architecture of the visual system reflects the constraint that two properties, such as *blue* and *yellow*, cannot co-occur at the same point, it is reasonable to infer that the *representation* as of a point as *blue* and the *representation* as of that point as *yellow* are structurally incompatible—the perceptual states cannot be combined to form a representation of a point as *blue and yellow*.

The structural interdependencies between perceptual representations of size and dis-

tance, shape and orientation, lightness/color and illumination, are rooted in the way the perceptual system must disentangle these pairs of properties. The perceptual system capitalizes on natural constraints in order to distinguish between distal non-relational properties such as shape and color and distal relational properties such as surface orientation and illumination, which have been confounded in the proximal stimulus. As I emphasized in the previous section, the architecture of the perceptual system reflects constraints on how those features must go together to produce the proximal stimulus. While I argued that we do not represent those constraints, they nevertheless seem to shape the structures of our perceptual representations, so that there are only certain ways of structuring representations of non-relational properties given how we represent the relevant relational properties. For example, the array-like structure of our perceptual representations of surfaces embodies the constraint that, given an elliptical pattern on the retina, you must be looking at either a slanted circle or a head-on ellipse.³⁹ An elliptical pattern on the retina must have been produced by a surface that occupies certain lines of sight. So, under normal conditions the representations of points on that surface will fill an elliptical pattern of cells in the representational array. Depending on how one represents the distances to (or depth relations between) the edges in those cells, the combination of the representations that fill those cells must represent either a slanted circle or a head-on ellipse.

The guiding idea behind the structural approach is that the perspectival character of representations depends on the structural features of those representations—on the nature of their representational ingredients and the ways those ingredients are combined. Perspectival

39. See Burge, “Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints,” p. 494–5.

variance and similarity result from differences and similarities in the way perceptual representations are structured from their parts. There is a common thread behind what makes different forms of perception perspectival. In each case, representations of non-relational properties such as size, shape, lightness, and color and representations of relational properties such as distance, orientation, and illumination are structurally interdependent. It is no accident that these structural interdependencies show up throughout perception. These interdependencies correspond to the way the perceptual system's architecture reflects natural constraints on how proximal stimuli confound distal properties. While the structural approach calls for more development, it promises to offer an empirically plausible, unified explanation of the perspectival character of perception.

VI CONCLUSION

Many have held that perception is perspectival because we perceive things as having certain relational properties that correspond to the varying ways those things look relative to different conditions of observation. The pluralist version of this view explains the perspectival character of spatial perception in terms of the perception of one kind of property, such as visual angle, while explaining the perspectival character of color perception in terms of another kind of property, such as surface luminance. This approach is unsatisfying because it fails to give a systematic, unified account of the perspectival character of perception. By contrast, the perspectival properties approach offers a unified account on which perception is perspectival because we always perceive things as having perspectival properties. This approach is untenable because it posits perceptual representations that would have no identifiable function in perceptual information processing.

I have proposed an entirely different explanation of the perspectival character of perception. Perceptual representations are perspectival because of the ways they are structured from their parts. What perspectival systems of representation have in common is that the structures of representations of non-relational properties (such as size, shape, lightness, and color) and representations of relational properties (such as distance, orientation, illumination) are interdependent. These structural interdependencies are rooted in the way perception works to disentangle the contributions that those properties make in producing sensory stimuli. Perspectival representations form an explanatorily unified psychological kind not by virtue of the properties that they represent but by virtue of how they are structured.