Analytic Philosophy Vol. 0 No. 0 June 2019 pp. 1–34

USE YOUR ILLUSION: SPATIAL FUNCTIONALISM, VISION SCIENCE, AND THE CASE AGAINST GLOBAL SKEPTICISM

E. J. GREEN
Massachusetts Institute of Technology

GABRIEL OAK RABIN
New York University, Abu Dhabi

1. Introduction

We perceive and think about spatial properties. But what is the nature of these capacities, and how do they relate to one another? *Spatial functionalism* is the view that spatial concepts, as well as spatial experiences, represent whatever worldly entities play a certain functional role. Our concept SQUARE could have turned out, in principle, to pick out rectangles, circles, or even thoughts in the head of Descartes' evil demon. Whether a precise spatial functionalist view sanctions such possibilities will depend on the details of the specified function. According to *causal phenomenal spatial functionalism*, our spatial concepts, as well as our spatial experiences, represent the normal causal bases of our spatial experiences. If squarish experiences are normally caused by rectangles, then our concept SQUARE refers to rectangles.

Chalmers (2012, 2019) defends causal phenomenal spatial functionalism: "Shapes such as squareness are picked out in virtue of their role in causing our experiences of shape. In fact, these cases suggest that our relevant concept of shape is a concept of whatever normally causes the relevant shape experiences" (2019: 21). He defends an analogous view of our concepts of size properties and left-right relations. The idea is that our spatial concepts pick out whatever properties normally cause the right sorts of spatial experiences. Searle (2015: ch. 4) takes a similar position, although the two views are not quite the same (see below).

One interesting implication of causal phenomenal spatial functionalism is that systematic, lifelong illusions with respect to spatial properties are impossible: "If we pick out spatial properties as the normal causes of our spatial experiences, then situations in which spatial experiences are normally caused by properties other than the spatial properties that they represent will be ruled out" (Chalmers 2019: 24). Let's label this claim the *no lifelong illusion thesis*.

No lifelong illusion thesis: Systematic lifelong illusions with respect to spatial properties are impossible.

Analytic Philosophy Vol. 0 No. 0

Chalmers argues that if the no lifelong illusion thesis is true, then certain allegedly global illusion scenarios, including brains in vats, Cartesian demons, and matrix scenarios are not illusory at all – at least as far as our spatial beliefs are concerned. The anti-skeptical conclusion applies to any brain-in-a-vat style scenario that meets two conditions (a) *life-long embedding*: the subject is embedded in the vat/matrix/demonic-dream-state for their entire life, and (b) *regularity*: the causes of the subject's experiences follow a regular pattern (Chalmers 2017). The regularity clause rules out chaos worlds in which purely random factors generate our experiences. It also guarantees that there is a normal cause of shape experiences. Causal phenomenal spatial functionalism explains how these scenarios could fail to be skeptical. If SQUARE refers to the normal cause of squarish experiences, and that cause is a certain program in the matrix, then SQUARE refers to that program. When inhabitants of the matrix say, "Lo, a square!", they are correct, not mistaken, because they are interacting with the program represented by their SQUARE concept.

Chalmers motivates the no lifelong illusion thesis primarily by arguing that our spatial concepts are twin-earthable.[1] He argues that on El Greco World, which is stretched 2:1 vertically but is otherwise a functional duplicate of the actual world (see Hurley 1998; Thompson 2010), our counterparts' squarish experiences represent 2:1 rectangles. Causal phenomenal functionalism makes this prediction, because 2:1 rectangles are the cause of El Greco squarish experiences. Chalmers takes these twin earth scenarios to suggest that certain types of systematic lifelong spatial illusions are impossible. He argues for causal phenomenal spatial functionalism on the grounds that it both entails and explains the no lifelong illusion thesis and the twin-earthability of spatial representations. In turn, the no lifelong illusion thesis is used to support an antiskeptical conclusion about our experiential and conceptual grasp of spatial properties.

Unfortunately, the no lifelong illusion thesis is false. Lifelong systematic spatial illusions are not only possible, they are actual. Evidence from vision science suggests that many of our spatial experiences are systematically nonveridical (McLaughlin 2016; Wagner 2006; Wagner and Gambino 2016). These illusions are not fringe cases involving bizarrely shaped rooms or challenging lighting conditions. They are commonplace and involve relatively ideal conditions.[2] This spells trouble for causal phenomenal spatial functionalism, and for Chalmers's case against skepticism.

[1] Roughly, a individual's concept C is *twin-earthable* if and only if one can devise a scenario in an intrinsic duplicate of the individual uses a counterpart concept to refer to a distinct property. For more on twin-earthability and El Greco world, see section 7 as well as Chalmers (2012), Rabin (ms).

[2] Lifelong illusion or error, in the sense of misrepresentation, does not entail lifelong sensory malfunction. It may be that when our perceptual systems are functioning properly, they misrepresent the environment in certain systematic ways. Thanks to an anonymous reviewer for pointing this out.

Still, we are sympathetic to a broadly functionalist theory of spatial representation. We are inclined to think that certain allegedly global illusion scenarios are not skeptical scenarios at all. But we also believe that contemporary vision science has demonstrated that we *homo sapiens* are subject to certain systematic lifelong illusions. This leads us to reject the version of spatial functionalism Chalmers proposes: causal phenomenal spatial functionalism. We examine ways in which the view could be modified in response and argue that most of the easy fixes do not work.

The structure of the paper is as follows. In Section 2, we discuss empirical evidence that the same shape property will normally cause distinct experiences depending on the orientation in which it is presented. Section 3 uses these empirical results to pose a dilemma for causal phenomenal spatial functionalism. Sections 4–6 argue that three potential routes out of this dilemma do not succeed. Section 7 considers the twin earth scenarios that Chalmers takes to support his view. We argue against the inference from the twin-earthability of spatial concepts to causal phenomenal functionalism. Section 8 considers the prospects for spatial functionalism to deliver an antiskeptical verdict about brains in vats, once we give up the causal-phenomenal version of the view. One surprising upshot is that if a version of spatial functionalism is true, then, as regards the veridicality of our spatial experiences, we could be better off in the matrix than we are in the real world.

Causal phenomenal spatial functionalism is a conjunction of two theses:

> **Causal phenomenal spatial experience functionalism:** For any spatial phenomenal experience-type E, E represents the worldly property that normally causes instances of E.
>
> **Causal phenomenal spatial concept functionalism:** For any spatial concept C, there is an associated experience-type E such that C represents the worldly property that normally causes instances of E.[3]

These two functionalist theses – concerning the reference of concepts and experiences – pair naturally.[4] The claim about the content of experiences is primary. Spatial concepts (SQUARE) inherit their

[3] We do not attempt to spell out, in any detail, what it means for an experience-type to be "associated" with a concept. The basic idea is clear. The experience-type associated with concept C is the type of experience on the basis of which we tend to apply C. Squarish experiences are associated with the concept SQUARE because we apply the concept SQUARE on the basis of squarish experiences. Reddish experiences are associated with RED for analogous reasons, etcetera. Chalmers (2019) is also committed to the idea that concepts and experiences can be associated in this way. Chalmers calls the associated experiences 'the relevant experiences' (21).

[4] We will occasionally speak of experiences *referring* to spatial properties. If the reader dislikes such talk, they should feel free to substitute whichever of 'represents', 'indicates', 'picks out', or 'has as its content' they prefer.

reference from associated experiences (squarish experiences). We apply spatial concepts on the basis of their associated experiences. We use the term 'phenomenal' to mark the fact that the relevant functional role is specified in terms of phenomenal experience. The qualifier 'causal' tells us what meta-semantic role the phenomenal plays: look for the normal cause of the phenomenal experiences.

Spatial functionalism can be construed as a thesis either about the semantic content or about the reference-fixing conditions of spatial concepts and experience — similarly for the causal phenomenal version of the view. Here, we primarily treat spatial functionalism in general, and causal phenomenal functionalism in particular, as meta-semantic theses about the factors that fix the reference of spatial concepts and experiences. Attacking the weaker reference-fixing thesis is sufficient to defeat Chalmers's endorsement of the content version. For Chalmers, content (primary intension) always involves reference-fixing conditions available to a priori reflection (Chalmers 2006a,b). If the "normal cause" condition does not fix reference, then it can't be knowable a priori or be a reference-fixing content.

Chalmers (2019) accepts both causal phenomenal functionalist theses. However, late in the paper he briefly conjectures that further factors, including holistic considerations, might also play a role in fixing reference for spatial representations. We first consider and argue against what is both the most straightforward and cleanest version of the causal-phenomenal functionalist theory and the most natural interpretation of Chalmers's statements of the view. (See also Stoljar (forthcoming) for the same interpretation.) That view is the conjunction of the two theses above. Section 5 will address the holistic strategy.

Chalmers (2017, 2019) holds that the causal-phenomenal reference-fixing analyses for our spatial concepts are knowable through a priori reflection. We can put this as the claim that for any spatial concept C, there is an a priori knowable truth of the form: 'C refers to the worldly property that normally causes instances of experience-type E'. We argue also against this claim of a-priority.

Searle (2015) joins Chalmers in proposing a broadly functionalist account of spatial concepts:

> [E]ven for shapes and lines there is a conceptual connection, a necessary connection, between the features of the object and its ability to cause certain sorts of experiences. Part of what it is for an object to be a straight line or to be a circle is to be able to cause this sort of experience.(Searle 2015: 128)

Searle, unlike Chalmers, does not argue for an antiskeptical interpretation of matrix-style scenarios. On Searle's view, causal relations to experience constitute a necessary, but not sufficient, condition on the reference of spatial concepts. The referent of SQUARE must normally cause squarish experiences. But the referent must also satisfy further

conditions, and programs of the matrix don't qualify (e.g., Searle 2015: 159-160). We won't examine Searle's account in detail, but our arguments challenge this view as well. For example, evidence suggests that squares do not normally cause squarish experiences, so this cannot be even a necessary condition on the reference of the concept SQUARE.

2. Depth Compression in Visual Space

The goal of this section is to establish two points. First, objects with the *same* external shape will normally cause *distinct* shape experiences depending on the orientation in which they are presented. Second, objects with *different* external shapes will normally cause the *same* shape experience depending on the orientation in which they are presented. Our evidence for this draws on work in vision science on the compression of visual space. The primary upshot of this literature is that we tend to perceptually experience objects oriented in depth as systematically contracted relative to objects oriented in a frontal plane. A *frontal plane* is any plane perpendicular to the perceiver's line of sight.

For city drivers, the following experience is commonplace. You are driving around looking for parking. You see a potential spot about a block away. From your perspective, however, the cars seem to be packed far too tightly to accommodate your vehicle. You approach and discover, to your surprise, that you can fit with room to spare. Happily, you back into the spot on your first attempt. This is a mundane observation, but it illustrates an important point: In normal, everyday contexts, our perception of length and distance is often compressed along the depth dimension.

Perceptual psychology provides evidence that the compression of perceived length in depth is *systematic*. The phenomenon is robust across stimuli, experimental task, instruction type, and viewing conditions. For example, in one typical paradigm, the subject views a scene containing a pair of fixed markers forming an imaginary line segment within the subject's frontal plane. The subject is then asked to adjust the position of a third marker that, together with one of the frontal markers, forms a segment oriented in depth (Figure 1). The subject's task is to adjust the depth segment until it appears equal in length to the frontal segment. The consistent result from experiments of this sort is that the depth markers are adjusted to form a larger interval than the frontal markers.[5] This holds regardless of whether subjects

[5] For this general pattern of results, see Wagner (1985), Loomis et al. (1992), Loomis, Philbeck, and Zahorik (2002), Tittle et al. (1995). Studies differ in the exact *degree* of depth compression, and in how the compression *changes* as a function of distance. The typical finding, however, is that compression grows more dramatic with increased distance from the subject. Thus, Loomis et al. (1992) found that at a distance of 4 meters, depth intervals appeared roughly ¾ the length of physically equivalent frontal intervals, while at 12 meters this fraction fell below ½.

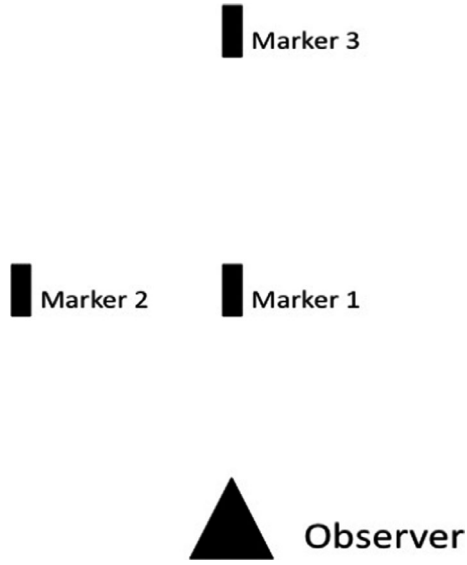

Figure 1: Schematic depiction of an experiment on depth compression. The observer is asked to adjust the interval between markers 1 and 3 until it appears equivalent in length to the interval between markers 1 and 2 (or vice versa.

are given *apparent* instructions (instructions to base their judgments on the way the stimulus subjectively *appears*) or *objective* instructions (instructions to base their judgments on the way they believe the stimulus *really is*).[6] A similar result is obtained when subjects are simply asked to verbally estimate the height-to-width ratio of an L-shaped object oriented in depth. They systematically underestimate length along the depth dimension (Loomis, Philbeck, and Zahorik 2002, exp. 2). The received explanation for these findings is that depth intervals appear shorter than physically equivalent frontal intervals. Thus, to make a depth interval perceptually appear the same length as a given frontal interval, the subject needs to adjust the former to be objectively longer than the latter.

Other studies have shown that the perception of angle is also affected by visual depth compression. Wagner (1985) showed that when the open side of an angle faces directly toward or directly away from a subject, the angle appears larger than a physically equivalent angle viewed from the side. When subjects need to adjust angles so that they visually 'match', angles viewed from the side are rendered

[6] In the case of apparent instructions, see the Wagner and Gambino (2016) study discussed below. In the case of objective instructions, see Loomis et al. (1992), Loomis, Philbeck, and Zahorik (2002).

physically larger than angles viewed straight on. As Wagner observes, this is readily explained by the depth compression account:

> [I]f angles seen facing directly toward or away from the observer are 'squeezed' in the in-depth dimension, they would perceptually expand in their size. (Points along the sides of the angle remain the same distance apart in the frontal dimension, but are moved closer to the vertex in the in-depth dimension. In the extreme case in which all in-depth distances approach zero, all points would be compressed along a single line, forming a 180° angle.)(1985: 486-488)

Studies have also found that the same general pattern (compression of depth intervals relative to frontal intervals) carries over to the perception of complex 3-D shape (Todd 2004; Todd et al. 2004).

While this is the basic lesson, one addendum should be noted. It turns out that depth intervals *very* near the subject (closer than about 1 meter away) are perceptually *expanded* relative to frontal intervals (e.g., Loomis, Philbeck, and Zahorik 2002; Tittle et al. 1995). Wagner and Gambino (2016) recently assessed perceived depth expansion and compression as a function of viewing distance. Under outdoor, full-cue conditions with binocular viewing, subjects saw a series of tape segments positioned either frontally or in depth, and they were asked to indicate the segment's size by adjusting a tape measure (with numeral markings hidden) until it appeared the same length as the tape segment they were evaluating. Subjects were given apparent instructions—they were told to indicate the length the segment looks or appears from their perspective (Wagner and Gambino 2016: 591). The authors found that lengths in depth were slightly expanded relative to frontal lengths at a distance of 1 meter from the viewer, but this quickly shifted to compression at distances beyond a meter. The relative depth compression reached an asymptote for distances beyond 7 meters, where in-depth intervals were perceived, on average, to be about half as long as physically equivalent frontal intervals the same distance away. For instance, a 1 meter interval oriented in depth would have been perceived, on average, as roughly equivalent to a 0.5 meter interval in the frontal plane. This pattern of expansion rapidly shifting to compression was also consistent with the results of a meta-analysis within the same article.

We take these results, together with the routine introspective observation noted above, to support the following conclusions. First, objects with *different* real-world shapes will normally cause the *same* type of visual shape experience if they occupy different orientations with respect to the perceiver. Consider, for instance, Chalmers's 'squarish' experience. The research just discussed indicates that this type of shape experience is caused by at least two different types of objects: Certain rectangles seen in the frontal plane at a distance of

10 meters, and certain rectangles of about *twice* the height-to-width ratio seen lying on the ground at the same distance. Likewise, objects with the *same* real-world shape normally cause *different* types of visual shape experience depending on their orientation with respect to the perceiver. A physical square 10 meters away causes one type of shape experience when seen in the frontal plane, but an experience that we would describe as an experience of a more compressed, less squarish, shape when seen lying on the ground.

3. A Dilemma for Causal Phenomenal Spatial Functionalism

Following Masrour (2017), we will refer to the experiences produced by a square viewed in the frontal plane and the same square viewed in depth as *seemingly incompatible*. They are seemingly incompatible because they seem to present distinct shape properties. Masrour defines seeming incompatibility as follows:

> Two experiences of the same object, E_1 and E_2, are seemingly incompatible with respect to a determinable quality, Q, when *in virtue of their phenomenal characters* (a) E_1 seems to present the object as having the property P, (b) E_2 seems to present the object as having the property R, and (c) the subject of experience is disposed to judge that P and R are incompatible values of Q.(2017: 568)

Let Q be the determinable *shape*. Perceptual psychology suggests that (a) the experience normally caused by a physical square oriented frontally, and (b) the experience normally caused by a physical square oriented in depth seem to present two distinct shape properties. Call the experience in (a) $E_{frontal}$, and call the experience in (b) E_{depth}. E_{depth} seems to present a shape that is compressed along one axis relative to the shape presented by $E_{frontal}$. Moreover, the subject is disposed to judge that these experientially presented shapes are incompatible values of the shape determinable. How do we know the latter fact? We know it because they *in fact* make this judgment (and make it systematically) across a variety of contexts. Thus, $E_{frontal}$ and E_{depth} are seemingly incompatible. They seem to present distinct and incompatible shape properties.

Critically, we are not assuming that either $E_{frontal}$ or E_{depth} is an experience as of the property *squareness*. (In fact, as we'll argue below, it is likely that *neither* is an experience as of squareness.) Our labels are intended to type-identify the experiences by their phenomenal character, while remaining explicitly noncommittal which properties the experiences represent. The important observation is just that $E_{frontal}$ and E_{depth} have distinct phenomenal characters, and in virtue of their distinct phenomenal characters they *seem* to present distinct shape properties.

We now argue that causal phenomenal functionalism faces a dilemma. We have two seemingly incompatible experiences: $E_{frontal}$ and E_{depth}. Either (i) the causal phenomenal functionalist can hold that at least one of the two experiences is nonveridical, or (ii) she can hold that both experiences are veridical. Either way, the view is in trouble. Both (i) and (ii) require us to give up the no lifelong illusion thesis. As such, both require us to reject the primary claim taken to support causal phenomenal functionalism.

Consider (i) first. This option requires that either $E_{frontal}$ or E_{depth} is nonveridical, but notice that the discrepancy between frontal and in-depth experiences of the same object is wholly general. It holds across stimuli and across perceptual contexts. Accordingly, the advocate of this approach will be forced to say that either frontal shape experiences or in-depth shape experiences are normally nonveridical. But then, contra the no lifelong illusion thesis, we are subject to at least one systematic, lifelong illusion. Either we systematically misperceive objects oriented in the frontal plane or we systematically misperceive objects oriented in depth. This would be a case of reliable misrepresentation – getting things wrong in the same way all or most of the time (Mendelovici 2013). Either perception systematically overestimates the lengths of intervals in the frontal plane or it systematically underestimates the lengths of intervals oriented in depth. But recall that Chalmers's argument for causal phenomenal functionalism relies on the claim that systematic, lifelong spatial illusions are impossible. If we accept (i), then not only do systematic, lifelong spatial illusions fail to be impossible — they are actual.

Now consider (ii). This option requires that both experiences are veridical, despite their seeming incompatibility. Presumably, then, both experiences represent the property *square*, but in different ways, or under distinct modes of presentation. This seems like the more promising route for the causal phenomenal functionalist because it appears to preserve the no lifelong illusion thesis. However, we believe that even on option (ii), we are forced to sanction another systematically illusory experience.

Suppose that a perceiver views two square objects, A and B, simultaneously. A is viewed straight on; B is oriented in depth. A elicits an instance of $E_{frontal}$; B elicits an instance of E_{depth}. These are constituents of the subject's overall perceptual experience of the scene.[7] Option (ii) requires that these experiences veridically represent the shapes of A and B, respectively. The perceiver is thus not subject to illusions about the *intrinsic* shapes of A and B. However, intrinsic shapes are not the only geometrical properties we perceive. We also perceive geometrical relations among objects. Objects can perceptually appear longer than, taller than, or wider than others. Accordingly,

[7] We assume here that experiences can have other experiences as parts (Bayne 2010; although see Tye 2007).

it seems that our perceiver is still subject to a further illusion. Her perceptual experience represents object B as shorter than object A along one of its axes, when in fact B is exactly as long as A along that axis. Again, because the phenomenon of depth compression is perfectly general, perceptual illusions about geometrical relations will be lifelong and systematic. Option (ii) allows the causal phenomenal functionalist to avoid the result that $E_{frontal}$ and E_{depth} systematically misrepresent their objects, but commits them to systematic illusions about the relations between these objects. Thus, they are still forced to reject the no lifelong illusion thesis.

It is plausible that perceptual experience represents relations like *longer than* and *shorter than*. However, the stalwart causal phenomenal functionalist might resort to denying this claim in order to preserve the no lifelong illusion thesis. They might suggest that our perceptual experience of geometrical properties is confined to the intrinsic shapes of objects. Geometrical relations between objects can only be reached through postperceptual inference. In the case of $E_{frontal}$ and E_{depth}, this postperceptual inference is faulty, leading to false beliefs about geometrical relations. Nonetheless, perceptual experience is perfectly veridical with respect to shape.

This picture of visual phenomenology strikes us as obviously wrong. But even if we grant it, the causal phenomenal functionalist faces further difficulties. First, she must explain why the subject is disposed to judge that object B is compressed relative to object A, given that her perceptual experience represents both objects as perfectly square. Second, the resulting picture still yields systematic lifelong mistakes in our perceptual *beliefs* about spatial properties. While our experience does not get the relative length relations between the two squares wrong (because we do not undergo experiences of these relations), our beliefs do. Recall that one of the main upshots of Chalmers's view was its ability to deliver an antiskeptical verdict about certain allegedly skeptical scenarios. But if the causal phenomenal functionalist adopts the strategy under consideration, they must grant that we all hold systematically mistaken beliefs about certain spatial properties. The antiskeptical result has been lost.

Third, the strategy presupposes that there is a sharp distinction between the perception of intrinsic shape and the perception of geometrical relations. The idea is that we perceive properties of the former type, but not the latter. But, plausibly, there is no sharp distinction between these abilities. Perceiving something as cubical is at least partly constituted by perceiving its sides as having the same shape (square). Accordingly, perception of intrinsic 3-D shapes poses the same threat to the no lifelong illusion thesis as perception of geometrical relations among 2-D shapes. Consider the perceptual experience of a 3-D cube viewed straight on at a distance of 10 meters. Because compression along the depth dimension is a systematic phenomenon, the perceptual experience of the object's top face will be

compressed relative to the perceptual experience of its front face.[8] Thus, assuming that perceptual experience represents the object's intrinsic 3-D shape, it will be experienced not as intrinsically cubical, but as a slightly compressed rectangular prism. This is an illusion about the object's *intrinsic* shape, not about its relations to other objects. Once again, the phenomenon of depth compression suggests that this type of illusion about 3-D shape is lifelong and systematic. If this is right, then to rescue the no lifelong illusion thesis, one will need to deny not only that we perceive geometrical relations among objects. One also needs to deny that we perceive certain intrinsic shape properties of objects – namely, their 3-D shape properties.[9]

We think that the no lifelong illusion thesis is false. This shows that the primary datum taken to motivate causal phenomenal functionalism is incorrect. However, even if the functionalist finds a way to preserve the no lifelong illusion thesis, she will need to meet a further challenge. For, even if it is true that both $E_{frontal}$ and E_{depth} are veridical, and so that both represent the property *square*, we doubt that it is a priori that both represent the property *square*. Recall that, according to Chalmers (2012, 2019), the truth of causal phenomenal functionalism is knowable by a priori reflection. As regards the concept SQUARE, then, there must be at least one a priori knowable truth of the form 'the concept SQUARE represents the worldly property that normally causes experiences of type E.' But *which* experience-type should we substitute for "E" in this schema? There are two possibilities.[10]

A Priori Frontal: It is a priori knowable that the concept SQUARE represents the property that normally causes $E_{frontal}$ experiences.
A Priori Depth: It is a priori knowable that the concept SQUARE represents the property that normally causes E_{depth} experiences.

[8] See Tittle et al. (1995) for evidence that depth compression affects the perception of intrinsic 3-D shape for cylindrical objects.

[9] It might be suggested that a parallel argument can be run for the experience of intrinsic 2-D shape. Consider the perceptual experience caused by a square oriented in depth. Due to depth compression, we will experience the figure's front and back sides as longer than its left and right sides. But it is also plausible that we experience something as a square in virtue of experiencing all four of its sides as having the same length. If so, then depth compression already suggests that we are under illusions about intrinsic 2-D shape. We think this is a fair point. However, one might suggest that certain intrinsic 2-D shapes are perceptually *primitive*. We might perceive a figure as square, but not in virtue of perceiving length relations among its sides. If so, then we experience squares oriented in depth as intrinsically square, despite perceptual compression. But it is far less plausible that 3-D volumes are perceived this way. It seems clear that we perceive 3-D polyhedra by virtue of perceiving relations among their 2-D faces. When you perceive a figure as a stretched-out rectangular prism, this is because you perceive some relative-size relations among its faces.

[10] There are, of course, many other possibilities besides these, corresponding to the wide range of distances and orientations at which we can perceptually encounter square objects.

We envisage two options for the causal phenomenal functionalist. She might claim that just one of these options holds, or she might claim that both of them do. We'll consider the first option, which effectively privileges one of the seemingly incompatible shape experiences over the other, in the next section. For now, we consider whether it could plausibly be a priori *both* that squares normally cause $E_{frontal}$ experiences *and* that squares normally cause E_{depth} experiences.

If both claims are knowable a priori, then it should also be knowable a priori that objects of the *same* shape (squares) normally cause both $E_{frontal}$ experiences and E_{depth} experiences. But this truth is not knowable a priori. This surprising fact was revealed to us by contemporary vision science. For all we know a priori, it is possible that the shape that normally causes E_{depth} experiences *really is* compressed relative to the shape that normally causes $E_{frontal}$ experiences, just as we are inclined to judge. But if this scenario, in which the objects are as just we judge them to be, is epistemically possible, then it is not a priori that $E_{frontal}$ experiences and E_{depth} experiences are normally caused by the same shape.[11] In fact, we think the situation is even worse. E_{depth} seems to present a shape that is more compressed than the shape presented by $E_{frontal}$. On the basis of a purely priori reflection, the *more likely* scenario is one in which the shape that causes E_{depth} really is compressed relative to the shape that causes $E_{frontal}$. A defender of the present strategy would need to claim that we can rule out this plausible scenario a priori.

This highlights an important point. Some of the problems we've raised for causal phenomenal functionalism generalize to causal theories of perceptual representation more generally. Any view on which perceptual experiences represent their normal cause has trouble with the phenomenon of seeming incompatibility.[12] After all, seemingly incompatible experiences are caused by the same shape property, but seem to present different shape properties. It is counterintuitive to claim that, despite appearances, such experiences in fact represent the same shape property. Chalmers's view, however, faces a more serious difficulty. For even if it is *possible* for seemingly incompatible experiences to represent the same external spatial property, it is surely not a priori that they do.

To sum up: The empirical phenomenon of depth compression poses a dilemma for causal phenomenal functionalism. Objects with

[11] This holds by Chalmers's own lights. According to Chalmers, if S is epistemically possible, then ~S is not a priori knowable (Chalmers 2012: 61). Thus, if it is epistemically possible that $E_{frontal}$ experiences and E_{depth} experiences are not caused by the same shape, then it cannot be a priori that they *are* caused by the same shape.

[12] For example, Masrour (2017) uses seeming incompatibility to pose a challenge for naturalistic representationalist theories of perceptual experience. We also agree with McLaughlin (2016) that naturalistic representationalists have trouble accommodating the normal misperception of spatial properties.

the same objective shape will normally cause distinct shape experiences depending on the orientation in which they are perceived. The functionalist needs to say which of these experiences is veridical. Either she can say that just one of the experiences is veridical, or she can say that both are veridical. Both options appear to sacrifice the no lifelong illusion thesis. The first option simply concedes that we are under certain lifelong illusions about intrinsic 2-D shape: Either we systematically misperceive the shapes of objects oriented frontally, or we systematically misperceive the shapes of objects oriented in depth. The second option avoids systematic illusions with respect to intrinsic 2-D shape, but still saddles us with systematic illusions with respect to both shape relations and intrinsic 3-D shape. Finally, even if causal phenomenal functionalists were to reject the claim that we perceptually experience geometrical relations and 3-D shapes, and thus preserve the no lifelong illusion thesis, it is still doubtful that the truth of causal phenomenal functionalism is knowable a priori, as Chalmers contends.

In Sections 4–6, we'll consider some routes out of this predicament that have some *prima facie* plausibility. But before doing so, we'll consider one more reply on behalf of the causal phenomenal functionalist. Consider the problem raised for option (ii) above. We argued that even if $E_{frontal}$ and E_{depth} are veridical as regards the intrinsic shapes of objects A and B, respectively, there is a further experience that is systematically illusory: A and B are experienced as having different intrinsic shapes, when in fact they are both squares. However, someone might deny that the experience of A and B as different in shape is illusory. Specifically, one might grant that we typically experience squares in the frontal plane as having a different shape from squares lying on the ground, but insist that such experiences are perfectly veridical.

The key is to deny that shape experience represents objective shape. Suppose that $E_{frontal}$ and E_{depth} instead represent *perspective-dependent* shapes. For example, it might be suggested that both experiences represent *P-shapes* in Noë's (2004) sense. An object's P-shape is the shape that it would project onto the subject's frontal plane. A circle seen straight on has a circular P-shape, while a circle seen at a slant has an oval P-shape.[13] Because these are distinct properties, there is no perceptual illusion involved in representing them as distinct. Thus, we can consistently hold (a) that $E_{frontal}$ and E_{depth} are both veridical, and (b) that the subject veridically experiences their objects (A and B) as differing in P-shape. This strategy might be extended to cover the perception of 3-D shape properties as well, if

[13] A similar option is to analyze perspective-dependent shapes in terms of the *solid visual angle* subtended by an object relative to a viewpoint (Green and Schellenberg 2018). We believe that the same problems arise regardless of whether perspective-dependent shapes are construed as P-shapes or solid visual angles. Thus, we do not distinguish the two options.

such properties can be construed in an perspective-dependent way. However, we won't go into detail about this because we believe that the proposal is both incorrect and inadequate to Chalmers's needs.

First, regardless of whether perceptual experience represents perspective-dependent shape properties, there is strong reason to think that perceptual experience *also* represents objective shape properties. And, as we saw above, it is the perception of objective shape properties that creates trouble for the no lifelong illusion thesis. As others have observed, the P-shape proposal overlooks the role of experienced depth in shape phenomenology. An oval viewed straight on shares its P-shape with a circle viewed at a slant, but the two are manifestly not experienced as having the same shape (Briscoe 2008). Furthermore, it is a central aim of our perceptual systems to recover objective, perspective-invariant shape properties. Perception of these properties is critical for applying recognitional concepts to objects (Biederman 1987). To recognize an object as a penny, you need to know that it is objectively circular, not that it presents a circular P-shape relative to your viewpoint. Contemporary models of human shape perception incorporate this insight (Pizlo 2008), and cognitive neuroscientists have discovered areas of high-level visual cortex responsive to perspective-invariant shape properties (Lescroart and Biederman 2013). Moreover, there is a thriving research program in vision science dedicated to the study of shape constancy: the recovery of invariant, objective shape properties across changes in proximal stimulation.

Second, even if we granted that perceptual experience represents P-shapes and not intrinsic shapes, this would still not be adequate to the causal phenomenal functionalist's needs. According to this view, the concept CIRCLE refers to the property that normally causes a certain type of shape experience. Call the relevant experience a 'circleish' experience. For the above reference-fixing analysis to work, the functionalist requires that circleish experiences are normally caused by circles *alone*, to the exclusion of other shapes. However, if the P-shape account is right, then this reference-fixing analysis cannot be correct. After all, if shape experiences are individuated in terms of the P-shapes that they represent, then the *very same* shape experience will be normally caused by *myriad* external shape properties. To see this, suppose that we construe a 'circleish' shape experience as an experience of a circular P-shape. Then a circleish shape experience will be normally caused *both* by circles viewed straight on *and* by elongated ovals viewed at a slant. Thus, it is false, on this view, that circleish experiences are normally caused by circles alone, to the exclusion of other shapes.

The upshot is this. Our concept CIRCLE simply is not the concept of whatever property normally causes circular P-shape experiences (or any other sort of P-shape experience). Our concept CIRCLE applies to pennies viewed straight on, but not to elongated ovals viewed at a

slant, even though both objects have a circular P-shape. Thus, even if perspective-dependent shapes can be used to rescue the no lifelong illusion datum (and, to be clear, we do not think that this is true), they cannot be used to rescue causal phenomenal functionalism.

4. Privilege the Frontal Plane

Sections 4–6 consider replies on behalf of the causal phenomenal functionalist. All three attempt to render the view compatible with certain systematic lifelong illusions. The first and third replies adopt the same basic strategy. They privilege a certain class of representations, applying the causal-phenomenal functional story to them, and take the reference of a wider class of representations to be derivative. This strategy will permit certain lifelong illusions, but only if those illusions are not in the privileged class.

Above we saw that squares oriented in depth and squares within the frontal plane give rise to seemingly incompatible experiences: $E_{frontal}$ and E_{depth}. This suggested that at least one of these experiences must be systematically nonveridical. If so, that experience cannot satisfy the causal phenomenal reference-fixing analysis. In reply, the functionalist might privilege one of the experiences – presumably $E_{frontal}$, the experience of squares in the frontal plane. According to this approach, $E_{frontal}$ experiences do represent the property (square) that normally causes them. Experiences of shapes oriented in depth inherit their content from associated experiences of shapes within the frontal plane. Thus, the experience as of a square oriented in depth (which, recall, is in fact normally caused by rectangles stretched along the depth dimension) inherits its content from experiences as of squares in the frontal plane. Call the experience as of a square oriented in depth $E_{depth\text{-}square}$. To emphasize, $E_{depth\text{-}square}$ is a distinct experience from E_{depth}. E_{depth} presents a compressed rectangle, not a square, but is normally caused by squares. $E_{depth\text{-}square}$ presents its objects as square, seems to present the same shape as $E_{frontal}$, and is normally caused by rectangles.

The privileging of frontal square experiences such as $E_{frontal}$ is not overly ad hoc. Privileging the plane orthogonal to line of sight is certainly more principled than, say, privileging the plane 10 degrees off orthogonal. It is also worth noting that when we desire the best, most accurate, perspective on an object, we typically shift our perspective to view the object in the frontal plane. If I want to know whether that serving dish is a circle or an ellipse, I hold it up so that the shape occupies the frontal plane.

Privileging frontal square experiences entails that systematic illusion for non-frontal square experiences is possible. This move renders the view compatible with the empirical finding that intervals extended in depth appear shorter than physically equivalent intervals in the

frontal plane. Experiences of shapes oriented in depth (such as E_{depth}) are nonveridical, while their frontal counterparts are veridical. But the view under consideration is not compatible with systematically illusory frontal square perception. Unfortunately, systematically illusory experience of frontally presented squares (i) cannot be ruled out a priori, as the causal phenomenal story maintains, and (ii) occurs in actual human perception.

Even if the true meta-semantic theory of human square perception turns out to privilege frontal perception, this privileging cannot be revealed through a priori reflection. Imagine a creature who perceives 2-D shapes in both the frontal plane as well as various planes extending in depth, but whose primary evolutionary goals – detecting food, avoiding predators, finding mates – are all served by accurate perception in the depth planes. The creature perceives shapes in the frontal plane, but these cases are ancillary fringe cases. If we are to privilege a perceptual plane in determining the reference of this creature's shape percepts and concepts, it would be perverse to choose the frontal plane. Doubly perverse to do so on a priori grounds, before knowing anything about the actual circumstances in which the being's perception operates. For all we know a priori, the imagined creatures are *homo sapiens*. For these reasons, we think it is a mistake to privilege frontal perception on a priori grounds.

It is also a mistake to privilege frontal perception on empirical grounds, or at least to privilege frontal perception in the way envisaged by our target causal phenomenal functionalist. Systematic frontal shape illusions are not only a priori possible. They actually occur in human vision. Above we relied on the fact that perceptual experience of depth intervals is compressed relative to the perceptual experience of frontal intervals. But a similar phenomenon (though less extreme) occurs within the frontal plane. Vertical intervals in the frontal plane are systematically seen as longer than physically equivalent horizontal intervals under normal viewing conditions (Hibbard et al. 2012; Wolfe, Maloney, and Tam 2005). This is known as the *vertical-horizontal illusion*.

Consider the perceptual experience of an upside-down "T" in the frontal plane (Figure 2). The experience of the horizontal line, $E_{horizontal}$, and the experience of the vertical line, $E_{vertical}$, seem to present lines of different length. Once again, the causal phenomenal functionalist has two options. Either one or both these experiences is nonveridical, or both are veridical. Either way, we are subject to systematic illusions with respect to lines and edges in the frontal plane. On the first option, perception either systematically overestimates vertical intervals or it systematically underestimates horizontal intervals. Thus, lines in one of these orientations are subject to normal misperception. The second option allows us to be veridical about the *objective* lengths of vertical and horizontal lines, but we are still systematically wrong about their *relative* lengths. Thus, the functionalist cannot stave

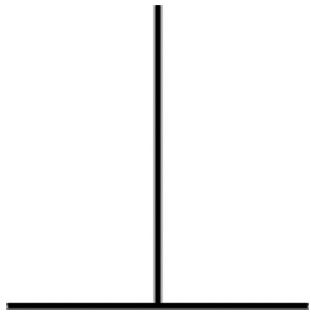

Figure 2: The vertical-horizontal (T-) illusion. ***Source:*** **Wolfe, Maloney, and Tam (2005).**

off systematic spatial illusion by privileging experiences of shapes in the frontal plane.

Moreover, note that while the vertical-horizontal illusion is extraordinarily simple, the misperception of relative length between vertical and horizontal lines has wide-reaching consequences for the perceptual experience of shape. As McLaughlin (2016) notes, "[C]ircles don't look circular to a normal perceiver in normal viewing circumstances. Rather, circles look slightly oval." In an elegant study, Hibbard et al. (2012) confirmed that this is true. Participants were shown a cluster of dots within an elliptical spatial region, giving rise to the percept of an ellipse with a border corresponding to the boundaries of the region. Binocular disparity cues[14] were provided through goggles to manipulate the perceived slant of the ellipse. These cues could either specify that the figure was located within the frontal plane or they could specify that it was slanted away from the frontal plane. The subject's task was simply to adjust the figure's aspect (height-to-width) ratio until it appeared circular. Strikingly, the experimenters found that there was *no* slant at which the figure perceived as circular filled an objectively circular region on the computer screen. Hibbard et al. conclude: "At all slants tested, the surface that appeared to have a circular outline was shorter than it was tall in the image" (41). Thus, the systematic misperception of relative length in the frontal plane leads to systematic illusions about aspect ratio, and this systematically infects the perception of 2-D shape.

5. Holism and Equivalence Classes

In most places, Chalmers (2019) seems to accept the most straightforward version of causal phenomenal functionalism, according to which

[14] Binocular disparity is, roughly, the amount of displacement between an object's location in the left and right retinal images (e.g., Palmer 1999: ch. 5). It is standard practice in vision science to manipulate perceived depth and 3-D orientation by altering the disparity between the images presented to the two eyes.

"our relevant concept of shape is a concept of whatever normally causes the relevant shape experiences." However, at certain points Chalmers backs off a bit, suggesting that if the reference-fixing "works in a holistic way," then certain lifelong illusions will be permitted (24). The holistic suggestion is difficult to assess because few details are provided, but we contend that the strategy is unlikely to succeed.

A natural implementation of the holistic strategy uses holistic considerations to divide experiences into equivalence classes (e.g. the class of all squarish experiences) and assigns the same reference to all members of that equivalence class. Specifically, Chalmers (p.c.) suggests that we collect into an equivalence class all the length experiences that seem, to the perceiver, to present the same length. All those experiences get assigned the same content: presumably, some worldly absolute length. Because the two sides of $E_{\text{depth-square}}$ seem to be the same length, they will be placed in the same equivalence class and represent the same length. This guarantees that $E_{\text{depth-square}}$ represents *square*. Since $E_{\text{depth-square}}$ is normally caused by oblong rectangles, we have systematic illusion that is compatible with this version of spatial functionalism. So far so good.

Recall that Chalmers argued for spatial functionalism precisely on the grounds that squarish experiences could have represented 2:1 rectangles (recall El Greco world). Unfortunately, the current suggestion deprives spatial functionalism of its ability to deliver this very verdict, thereby undercutting a main motivation for the theory. El Greco squarish experiences, which are indistinguishable from ours, present the figure's horizontal and vertical sides as the same length. On the view under consideration, "same length as" appearances determine membership in an equivalence class, all of whose members are assigned the same content. The two El Greco length experiences will be placed in the same equivalence class and receive the same content (a worldly distance). This will make it impossible for one to represent a distance twice the other, which is exactly what representation of 2:1 rectangles by El Greco squarish experiences requires. But Chalmers argues for spatial functionalism precisely on the grounds that squarish experiences could have represented 2:1 rectangles.

We have a second, more general, concern about the strategy of assigning contents holistically via equivalence classes. Consider how the option might work in the case of squarish experience. The idea is that although certain (perhaps most) squarish experiences are systematically illusory, squarish experiences as a whole have a single most common cause. Squarish experiences in response to *this* property are guaranteed to be veridical. Let's assume, for simplicity and specificity, that this property is in fact *square*.[15] Crucially, the most common

[15] It is worth noting that there is no guarantee that the most common cause of squarish experiences is in fact squares. That it is an empirical wager. It is possible, given vision science, that the most common cause of your squarish experience is an oblong rectangle.

cause of any perceiver's squarish experience depends on vicissitudes of personal history. As a result, there are cases where the holistic strategy seems to make the wrong prediction. Consider Axl. Axl is a human perceiver with a visual system like ours, subject to the same psychophysical laws that we are. However, almost all of Axl's squarish experiences are caused by objects in the frontal plane. He just hasn't come across many objects that produce squarish experiences at other orientations. For all we know, there are perceivers like Axl in the actual world. Problem: The most common cause of Axl's squarish experiences will be *non*-squares (horizontally elongated rectangles). Chalmers thus needs to say that Axl's squarish experiences represent non-squares. But this seems clearly to be the wrong result. If Axl's visual system works the same way ours does, and he also shares both our environment and our evolutionary history, his spatial experiences should represent the same properties as ours. A random gap in his personal history of spatial experience surely should not make this huge a difference to the contents of his spatial experiences.

This objection generalizes far beyond the case of squarish experience. The point of the equivalence class strategy is to permit systematic illusion by allowing some members of the equivalence class of experiences to be systematically caused by a property that is not the most common cause of the equivalence class as a whole. Consider any such equivalence class of spatial experiences, some of whose members are subject to lifelong illusion. We contend that it is empirically possible for there to be a perceiver like us, with our evolutionary history, in our sort of environment, for whom the illusory experiences predominate.

The holistic strategy remains underdeveloped. There is certainly residual room for maneuvers. Equivalence classes are not the only way to incorporate holistic factors. However, we are willing to wager that any viable form of holism in this area must incorporate factors beyond causal relations between external properties and phenomenal experiences. And we warn against the danger of using holistic considerations as a panacea to ward off all potential counterexamples. Any appeal to holistic considerations should be supplemented with specific proposals.

6. Coping with Compositionality

6.1. Compositionality and Atomic Representations

To motivate the second strategy we will suggest on behalf of the causal phenomenal functionalist, we start with an objection. According to the simple version of causal phenomenal functionalism, every experience (or concept) refers to the normal cause of that (associated)

experience. Bozo the baboon's elliptical experiences (experiences as of an ellipse) refer to the normal cause, in Bozo's environment, of elliptical experiences. His horizontally elongated experiences (experiences as of objects wider than they are tall) refer to the normal cause of horizontally elongated experiences.

Now consider a horizontal elliptical experience, i.e. an experience as of an ellipse wider than it is tall. Assume that the normal cause of Bozo's horizontal elliptical experiences is neither elliptical nor horizontal. More precisely, the normal cause of his horizontal elliptical experiences has neither the property that is the normal cause of horizontally elongated experiences, nor the property that is the normal cause of elliptical experiences. The scenario is not hard to imagine. If Bozo has horizontal elliptical experiences only when he ventures into a cave full of light-refracting gems that make vertically elongated objects look horizontally elongated, and rectangles look elliptical, then the normal (and only) cause of Bozo's horizontal elliptical experiences will be vertical rectangles. Consider Bozo's horizontal elliptical experience. A horizontal elliptical experience is also a horizontally elongated experience. This horizontally elongated experience is not veridical; the object is taller than it is wide. Similarly, his elliptical experience is not veridical; the object is rectangular, not elliptical. Bozo seems to be under a systematic illusion. All of his horizontal elliptical experiences are nonveridical.

The basic problem is one of compositionality. Causal phenomenal functionalism says that an experience (or concept) refers to the normal causal source of that experience. In general, reference composes: 'silver didgeridoo' refers to the intersection of 'silver' and 'didgeridoo'. But "normal causal source of" does not compose. The normal causal source of horizontally elongated elliptical experiences need not be in the extension of 'normal causal source of horizontally elongated experiences' or in the extension of 'normal causal source of elliptical experiences'. This leads to problems for Bozo. Causal phenomenal functionalism entails that the referent (i.e. normal cause) of Bozo's horizontal elliptical experiences is not the intersection of the referents of his horizontal and his elliptical experiences.[16] The account predicts a veridical whole experience (veridical because Bozo's horizontal elliptical experience, which represents *vertical rectangle*, attributes this very property to a vertical rectangle) composed of falsidical parts: an experience as of horizontal elongation and an experience as of an ellipse. But, plausibly, if a complex experience's parts are not veridical, neither is the whole.

This compositionality challenge provides an a priori argument against causal phenomenal functionalism that is independent of the

[16] Note that advocates of causal-covariational psychosemantic theories have tended to restrict the causal requirement to primitives, leaving the contents of complex representations to be determined by compositionality (e.g., Fodor 1990).

considerations brought to bear so far. We think it highly problematic for a flat-footed formulation of causal phenomenal functionalism. A *flat-footed* approach applies the causal reference-fixing story to all spatial experiences and all spatial concepts. Thankfully, the objection points the way toward a solution. The obvious move, which we consider as the second strategy on behalf of the causal phenomenal functionalist, privileges primitive representations. It applies the functionalist reference-fixing story to those representations, and then derives the meaning of complex representations from the primitives. This story assumes that spatial representations, both in perception and in thought, are compositional. Complex spatial representations are built up from primitive representations. Suppose that our concept SQUARE just is the concept of a figure with four equal sides meeting at four right angles. Then the concept SQUARE can be analyzed in terms of more primitive geometrical concepts.

Determining the primitive elements of spatial experience is a difficult project. We do not attempt the formidable task of spelling out all the candidates. Instead, we lay out two conditions that the primitive elements must satisfy. Our preliminary investigation leads to a modest pessimism about the strategy of running the causal phenomenal functionalist story on primitive representations.

6.2. Can Spatial Primitives Save Causal Phenomenal Functionalism?

Fortunately for the causal phenomenal functionalist, it is plausible that spatial representations are compositionally structured. Geometry textbooks are shot through with definitions of geometrical terms in terms of more primitive terms. It is plausible that our perceptual and conceptual representations of spatial properties are similarly structured.[17] We are not skeptical of the idea that many perceptual and conceptual spatial representations are complexes composed from more primitive representations. Nevertheless, we are skeptical that they are composed from primitive representations that *also* meet the causal phenomenal reference-fixing criteria. And we are even more skeptical that an approach of this sort will vindicate Chalmers's antiskeptical conclusion about our grasp of spatial properties.

For the sake of simplicity, we restrict our attention to whether the compositional approach can be made to work in the case of spatial experience. However, we believe that essentially the same issues arise for spatial concepts. We'll also assume that experiential representations are genuinely compositional: complex experiences can be

[17] Considerations of productivity strongly favor a compositional approach. We can perceive (and can learn to recognize) a potentially limitless variety of shape properties. It is hard to see how the visual system could accomplish this impressive feat unless it represented complex shapes in terms of more primitive shape properties together with their spatial arrangement (see Biederman 1987; Hummel 2013).

generated from more primitive ones, where the content of a complex experience is a function of the contents of its parts. This claim might be rejected, but this is not our concern. If spatial experiences are not compositionally structured, then the compositional strategy is simply unavailable to the causal phenomenal functionalist.

To run the compositional story, two conditions must be satisfied. First, the functionalist must identify a set of primitive spatial experiences that provides sufficient representational resources to compose the rest of our spatial experiences. Thus, if experiential representations of *square* are not in the class of primitives, it should be possible to generate them by combining primitive experiential representations. Second, they need to argue that every primitive spatial experience S represents the property that normally causes S. This second constraint entails that lifelong systematic illusion at the level of primitive spatial experience is impossible. Unfortunately, it is very difficult to see how the causal phenomenal functionalist can meet both of these conditions together.

What types of experiences might meet the second constraint and thus be immune from systematic illusion? We have already seen that perceivers are subject to normal misperception of ordinary 2-D shapes like square and circle. This is true even of squares and circles presented in the frontal plane. Thus, experiences of ordinary 2-D shapes (let alone 3-D volumes) are not good candidates for inclusion in the class of primitive experiences. We must look to simpler experiences. There may be some highly generic spatial properties and relations that are immune from systematic illusion. For instance, perhaps experiences of topological properties like continuity and closure are systematically veridical. We know of no evidence to suggest otherwise. However, experiential representations of generic topological properties are not sufficient to generate experiential representations of ordinary shape properties. Circles, rectangles, and chiliagons are topologically equivalent, so perception could not "build" experiences of these properties merely by combining topological experiences. The problem, then, is to expand the list of primitives without bringing in representations that are subject to systematic illusion, given the evidence we've already discussed.

To appreciate the difficulty, consider a potential compositional analysis of squarish experience. Perhaps experiences of square are composed from experiences of *edge*, *angle*, and the relation of *congruence* or *equality*. Thus, we might experience a figure as square by experiencing it as a figure with four congruent sides and four congruent angles. The problem is that a crucial building block in this analysis is the experience of *congruence*. But this is precisely the sort of experience that perceptual psychology tells us is susceptible to widespread illusion. Perceptual experience makes systematic errors about which edges and angles in the world are equal to one another, and often these errors are quite significant. For a vertical edge to be

experienced as congruent with a horizontal edge, the horizontal edge must be objectively longer. And for an angle facing toward the perceiver to be experienced as congruent with an angle viewed from the side, the angle viewed from the side needs to be objectively larger.

This problem runs deep. The causal phenomenal functionalist needs a class of primitives sufficient for building up all of our experiences of geometrical properties. Accordingly, their task is rather like the task of axiomatizing a geometry. In the latter case, it is standard to start with some very basic elements (*point*, *line*, *betweenness*, etc.) and to define others in terms of them. The geometry is given by axioms that set constraints on how the primitive elements relate to one another. For instance, one of Hilbert's axioms states that for every pair of points, there is at most one line that contains them both. From the primitives and the axioms, one can construct analyses of complex 2-D and 3-D geometrical properties, like *rectangle*, *cylinder*, *ellipse*, and *chiliagon*. The goal for the causal phenomenal functionalist is to build up spatial experiences in an analogous way.

We now arrive at the core issue. It is plausible that spatial experience can represent Euclidean shapes of essentially arbitrary detail and complexity. (Consider, for instance, the experience of a highly detailed sculpture with many cavities and undulations.) The most natural strategy for the causal phenomenal functionalist, then, would be to base their analysis of experiential primitives on some known axiomatization of Euclidean geometry. After all, these axiomatizations offer the only sets of primitives that we *know* are sufficient to define Euclidean shapes of arbitrary detail and complexity. Problematically, however, congruence's foundational role goes beyond the toy analysis of squarish experience above. It is also ineliminable in Hilbert's axiomatization of Euclidean geometry (e.g., Blumenthal 1961: 51-53). Because experiences of congruence can't be included in the class of primitives, the functionalist cannot base their analysis on Hilbert's axiomatization.[18] Of course, Hilbert's approach is not the only option. Alternative models of Euclidean geometry instead treat *distance* as the sole primitive relation among points (e.g., Blumenthal 1961: ch. 7). But the functionalist fares no better following this route. Sections 3–4 argued that we are subject to systematic misperception of distance (either in the frontal plane, the depth dimension, or both). We leave it to the causal phenomenal functionalist to choose primitives that are both rich enough to generate all our variegated complex spatial experiences and simultaneously immune from systematic illusion. The foregoing considerations illustrate how difficult this task is. The functionalist had better choose carefully.

[18] Hilbert's axiomatization also includes the notion *straight line* as a primitive. However, there is evidence that our perception of collinearity or straightness is also subject to systematic distortion. Objectively straight lines standardly appear curved, while certain objectively curved lines appear straight (Cuijpers, Kappers, and Koenderink 2002; Todd et al. 2001).

We are doubtful that this project will succeed. But even if it did, the payoff would be meager relative to Chalmers's ambitions. Recall that Chalmers wishes to use causal phenomenal functionalism to argue that allegedly skeptical scenarios involving deviant normal causes for our spatial experiences do not involve systematic deception. This argument crucially relies on a general form of the no lifelong illusion thesis – one that applies to the majority of our spatial experiences and spatial judgments. For instance, Chalmers (2019) writes:

> The intuition that Cartesian skeptical scenarios – evil genius cases, Matrix cases, and so on – involve systematic deception turns largely on the intuition that they involve spatial illusions. If a subject has veridical spatial experiences and true spatial beliefs, they can hardly be said to be systematically deceived about the external world. And if lifelong spatial error is hard to sustain, lifelong skeptical scenarios are hard to sustain too. (24)

Elsewhere, Chalmers (2012: 335). The problem is that the view now under consideration, which restricts the normal-cause requirement to experiential primitives, *is* compatible with systematic, lifelong errors about spatial properties. The vast majority of our spatial beliefs concern complex shapes, not primitives like *edge*, *congruent*, and *continuous*. Complex shapes are critical for object recognition and categorization. But if the only spatial experiences immune from lifelong illusion concern basic geometrical elements and relations, then it remains possible that we are systematically mistaken about the spatial properties that *matter most* to us – the complex shapes of ordinary objects.[19] This possibility of systematic illusion applies to both the actual world and to Chalmers's target Cartesian-demon style scenarios. Thus, if causal phenomenal functionalists try to save the view through the compositional strategy, the antiskeptical payoff is lost.

7. Return to Twin-Earth

We have not yet considered Chalmers's primary argument in support of the no lifelong illusion thesis and, in turn, causal phenomenal functionalism. That argument rests on the claims that spatial concepts are twin-earthable and that twin-earthability supports the impossibility of lifelong illusion. Causal phenomenal functionalism, Chalmers maintains, provides the best explanation of these facts. This section addresses these twin-earth considerations. We agree with Chalmers that spatial concepts are twin-earthable. But we reject his move from

[19] Suppose, for instance, that primitive spatial experiences are mistaken only 5% of the time, and that errors among the primitives are uncorrelated. Then, despite a lack of systematic illusion at the level of the primitives, we would expect complex experiences composed of 14 or more primitives to be errant more than half the time.

this claim to causal phenomenal functionalism and the conclusion that lifelong illusions are impossible.

A *twin-earth case* for a spatial concept requires a pair of conceivable situations in which one subject's spatial concept represents one property, while the agent's intrinsic duplicate in the second situation uses a counterpart spatial concept to represent a different property.[20] Twin-earth cases are easy to cook up for certain concepts – WATER, GOLD, JULIUS – and much more difficult, if not impossible for others – BACHELOR, FRIEND, JUSTICE. Spatial concepts are often taken to fall into the first "non-twin-earthable" category (Chalmers 2010: 554, fn. 2). However, Chalmers (2019; see also Thompson 2010) disagrees.

Chalmers takes the El Greco world to motivate the twin-earthability of shape concepts. Recall that in El Greco world, everything is stretched vertically by a factor of 2, and squarish experiences are normally caused by 2:1 rectangles. Chalmers argues that our counterparts' spatial experiences are veridical. Their squarish experiences represent the property of being a 2:1 rectangle. The associated concept (which they express using 'square') represents the property of being a 2:1 rectangle.[21] Causal phenomenal functionalism vindicates these twin-earth intuitions.

We share Chalmers's intuition that the veridicality conditions for our counterparts' spatial experiences are different from ours. If our world had turned out like El Greco world, our SQUARE concept would have represented *2:1 rectangle*. Squarish experiences would have represented the same property. These intuitions provide compelling support for a functionalist approach to spatial representation. However, we do not accept that our counterparts' spatial experiences are immune from lifelong illusion. Furthermore, we disagree with Chalmers's claim that our El Greco counterparts' squarish experiences represent 2:1 rectangles because they represent their normal cause. If El Greco world perfectly duplicates the causal/functional structure of our world, then it duplicates actual psychophysics, and our counterparts are mistaken in just the same ways we are. Our squarish experiences are normally caused by objects in the frontal plane that are slightly wider than they are tall – say, 1:1.1 rectangles. In El Greco world, the corresponding experiences are normally caused by, say, 1.8:1 rectangles, but represent the property of being a 2:1 rectangle. They are subject to systematic illusion, just as we are.

[20] For more precise formulations of twin-earthability, see Chalmers (2012) and Rabin (ms).

[21] Chalmers extends Thompson's argument by proposing a physically possible analogue of El Greco world, which he calls "Lorentz earth". Chalmers argues that our Lorentz squarish experiences veridically represent 2:1 rectangles. (See Epstein (2018) for an in-depth discussion of the puzzles for shape perception raised by special relativity theory.)

One need not assume the no lifelong illusion thesis to permit the twin-earthability of spatial concepts. In principle, any version of spatial functionalism, as long as the functional role is satisfiable by properties other than the actual role-player, will permit the twin-earthability of spatial concepts. This point, together with the fact that we are subject to lifelong illusion, defeats Chalmers's claim that the causal phenomenal breed of spatial functionalism offers the best explanation of twin-earthability. This is all for the better, because we accept that shape concepts can be twin-earthed but we reject the no lifelong illusion thesis. Twin-earth considerations provide support for some version of spatial functionalism, but the move from twin-earthability to causal phenomenal functionalism is too hasty.

8. Prospects for Spatial Functionalism and Antiskepticism

8.1. Whither Spatial Functionalism?

We have argued against a certain version of spatial functionalism, causal phenomenal spatial functionalism, according to which our spatial concepts and experiences represent normal causes. We now turn to the prospects for spatial functionalism *simpliciter*, once we give up the causal-phenomenal version, as well as the prospects for spatial functionalism to deliver antiskeptical results about Matrix scenarios, brains in vats, and the like.

By rendering systematic spatial illusions impossible, causal phenomenal functionalism achieved a philosophical payoff. Certain allegedly skeptical scenarios turned out not to be skeptical after all. Could it turn out that an evil demon is playing tricks on us, and everything is twice as tall as we thought? "No," says the causal phenomenal functionalist. Whether one regards this consequence as a benefit will depend on one's other commitments. If one thinks an evil demon could trick us in this way (systematically, over our entire lifetime), then this result counts against causal phenomenal functionalism. If one thinks our experiences are veridical despite the demon's machinations, then this result counts in favor of the view. We find some version of the antiskeptical result appealing. We think that many allegedly skeptical scenarios are not, in fact, nearly as skeptical as many take them to be.[22] We also believe that many spatial concepts are twin-earthable. Without causal phenomenal functionalism in the stable to deliver these verdicts, and given our commitment to widespread spatial illusion, this leaves us in a prickly situation.

[22] Here, we are fundamentally in agreement with Chalmers (2005), who construes Matrix and brain in a vat scenarios are simply alternative theories about the fundamental nature of the world, i.e. rival "fundamental physics".

The way out is to find an alternative version of spatial functionalism. Chalmers uses 'spatial functionalism' as a label for what we have called 'causal phenomenal spatial functionalism'. But the basic idea of spatial functionalism is just that our spatial representations get their reference secured via functional criteria. Functional criteria, in general, are agnostic about the nature of the player of the functional role. Of course, the representatum must be able to play the role. But beyond that, all bets are off. As long as the computer algorithms that run the Matrix can play the functional role specified by the account, they can be picked out by our spatial representations.

Chalmers (2019) does the spatial functionalist a disservice by presenting his own view – causal phenomenal spatial functionalism – as the only functionalist game in town. He offers up presentationalism, a view according to which experiences directly present the nature of their referent, as the main option for someone who rejects his own view. He also argues that, if one rejects presentationalism, causal phenomenal spatial functionalism is "left on the table as the most plausible view" (9). Noncausal and nonphenomenal versions of functionalism are not considered.[23]

We have argued against the causal phenomenal species of the spatial functionalist genus. That species of view is, put bluntly, too simple. Representations do not get their meaning determined one-by-one, by their normal causes. Other factors, including compositional considerations, need to play a role. We have tried to point the way for how a successor view, which retains some of the spirit of causal phenomenal functionalism, could go. But such views face considerable obstacles (see Sections 4–6). We are doubtful that they can be overcome. The other functionalist strategy is to expand the functional role beyond causal relations to experience. This is the option we prefer.

There is something correct, we think, about the idea that causal relations between the perceiver and their environment will play a role in the true theory of what experiences represent. But this story is bound to be highly complex. Our spatial experiences do not simply represent whatever property normally causes them. The connection between concepts and the causes of their associated experiences is more tenuous. First, there is the simple point that by stepping back from the experiences themselves to the concepts, we introduce an extra link, and complication, in the chain. Second, the conceptual level leaves more room for broader theoretical considerations to play a reference-fixing role. For example, the reference of SQUARE, or

[23] Elsewhere, Chalmers (2017: 29-30) suggests two further constraints on the reference-fixing conditions for spatial concepts: (i) spatial properties tend to change continuously, and (ii) objects close together tend to have stronger interactions with each other. Although we doubt that these criteria alone are sufficient to determine the reference of our spatial concepts, perhaps they will figure in the true meta-semantic account.

LINE, might be partly determined by the roles those concepts play in geometrical theory, which could have little to do with the normal causes of certain experiences.[24]

Our concept SQUARE is not constrained to refer to whatever property normally causes squarish experiences. Rather, we suspect its reference is fixed by some combination of relations to other geometrical concepts, beliefs we have about how these concepts apply to things in our environment, interaction with perceptual states, and relations (including causal relations) to worldly properties. The true meta-semantics for SQUARE is considerably more complex than: 'SQUARE applies to whatever is experienced a certain way.' The relevant meta-semantic principles might involve measurement under ideal conditions, dispositions to roll down hills, etc. But they are likely to be functional in character. Things fall under the concept SQUARE because they play a certain role.

It is worth noting that there are extant psychosemantic theories that permit widespread error in the case of perception. For instance, teleosemantic theories permit a representation's content to come apart from its dominant cause (Millikan 1989). Spatial functionalists might take inspiration. The phenomenon of systematic error poses a puzzle for the spatial functionalist, not an insurmountable problem. The puzzle is to identify functionalist-friendly reference-fixing conditions that permit the sorts of systematic illusion that vision science has revealed.[25]

8.2. Whither Antiskepticism?

We now consider the prospects for antiskepticism once we give up causal phenomenal functionalism. We start by dividing skeptical scenarios into three classes. In a *lightweight skeptical scenario*, our perceptual abilities, recognitional capacities, or beliefs about certain spatial properties are systematically errant. In such scenarios, we can be largely mistaken about which particular objects are square, or circular, at both the level of perception and thought. A *middleweight skeptical scenario* is one in which *all* of our particular spatial experiences and spatial beliefs are false. Objects never have the spatial properties we take them to. To illustrate the difference, consider a scenario in which frontal-squarish experiences are veridical, but experiences of squares in other orientations are not. In such a situation, most of our

[24] Functionalist views of geometrical concepts are popular among mathematicians, including, most famously, David Hilbert. Shapiro (2000) writes that "Hilbert provided... *functional* definitions of terms like 'point', 'line', and 'plane'" (155). "Anything at all can play the role of the undefined primitives of points, lines, planes, and so on" (151).

[25] For an interesting attempt to accommodate normal misperception within the teleosemantic framework, see Ganson (2018).

squarish experiences are falsidical and we are systematically mistaken. But we have some particular spatial experiences that are correct: namely, (most of) the frontal squarish experiences. Middleweight is stronger than lightweight because a middleweight scenario requires our experiences and particular spatial beliefs to be always mistaken, whereas a lightweight scenario only requires them to be systematically mistaken. Conversely, a lightweight scenario is compatible with our experiences and particular spatial beliefs sometimes being correct. A middleweight scenario is not. Middleweight entails lightweight, but not vice versa.

Middleweight scenarios are compatible with the truth of certain spatial beliefs, such as THERE ARE SQUARES, which we call *non-particular*. A *particular* spatial experience/belief is an experience/belief about the positive spatial features of some particular object in our environment. Non-particular spatial beliefs will involve quantifiers, negation, comparisons between properties, or the like. Middleweight scenarios are compatible with the truth of non-particular spatial beliefs. We can be correct in thinking SQUARES EXIST or THAT IS NOT A SQUARE.

In a *heavyweight* skeptical scenario, our spatial concepts and experiences either fail to represent any properties at all, or represent properties that no worldly object instantiates. Nothing falls under the extension of our concept CIRCLE. Some have held that color experiences fall prey to heavyweight skepticism (Pautz 2009). We experience objects as blue, but no such property is instantiated in our world. A heavyweight skeptical scenario for CIRCLE is more difficult to imagine. The matrix is commonly interpreted as both a light- and middleweight skeptical scenario. The matrix inhabitants' beliefs about which of the particulars they encounter are square are all false.[26] All they encounter are computer programs, which do not have any physical shape. However, even if the matrix simulation itself does not contain any squares, the world as envisioned in the film does: some of the servers constructed by the AI overlords who run the matrix are square. If this common interpretation were correct, then the concept SQUARE would have an extension in the world of the matrix. Some of the matrix inhabitants' non-particular spatial beliefs, including SQUARES EXIST, are true. This matrix world fails to provide a heavyweight skeptical scenario. However, we can rejigger the example. Imagine the 1-atrix, which is exactly like the matrix with one crucial difference: the 1-atrix simulation is run by AI overlords who inhabit a one-dimensional universe.[27] If we suppose that the 1-atrix-inhabitants' spatial concepts do not pick out computer programs, and stipulate that

[26] To emphasize: This is the *common* interpretation of the matrix. It is not the one we (or Chalmers) accept.

[27] For those skeptical that this makes sense, it is worth recalling an interesting result: a one-dimensional computer (Turing machine) can compute anything a multi-dimensional computer can. (Lewis and Papadimitriou 1981).

the "real world" that hosts the 1-atrix simulation is 1-dimensional, then SQUARES EXIST is false. We would have a heavyweight skeptical scenario, and any heavyweight skeptical scenario is also both middle- and light-weight.

We have defined the three types of skeptical scenario in terms of all spatial experiences and concepts. Many more distinctions could be made. One could restrict the scenarios to only experiences or only concepts. One could restrict via particular concepts/experiences (SQUARE) or classes thereof (e.g. polygons). Or one could expand to cover nonspatial concepts. (We already mentioned one view (Pautz (2009)) according to which the actual world is a heavyweight skeptical scenario with respect to color experience). We leave the task of further dividing logical space for another time, and lump all spatial representations, experiential and conceptual, together and treat matters as all or none.

Causal phenomenal spatial functionalism entails that as long as a situation meets the two conditions of *life-long embedding* and *regularity* of cause, it cannot be a skeptical scenario of any weight. In effect, Chalmers argues that these heavy- and middle- weight scenarios are impossible precisely because lightweight scenarios that meet the two conditions are impossible. We disagree. *Contra* Chalmers's no lifelong illusion thesis, there are situations that satisfy both conditions but remain skeptical scenarios. First and foremost, assuming that we inhabitants of actuality have not recently been embedded in a matrix, contemporary vision science has shown that the actual world provides a counterexample: it is a lightweight skeptical scenarios for many of our spatial experiences (e.g. squarish and same-length-as experiences).[28] Genuinely skeptical Matrix/brain-in-a-vat/evil-demon scenarios satisfying *life-long embedding* and *regularity* can also be devised. We are inclined to think that The Matrix, as envisioned in the film, has broadly the same psychophysics as the actual world, and thus is a lightweight skeptical scenario, just like the actual world. We leave it an open question whether there are middle- and heavy-weight skeptical scenarios that meet Chalmers's two conditions. When it comes to perceptual illusion, much depends on the particular psychophysics of the scenario in question, and not only on whether the perceiver has been embedded in the scenario for their whole life and whether the causal sources are regular rather than chaotic. All of the foregoing spells trouble for causal phenomenal functionalism. But Chalmers's strategy of ruling out middle- and heavy-weight scenarios by ruling out lightweight scenarios is like picking up the desk to move your computer. There are other, simpler ways.

[28] Spatial beliefs are another matter, because we have alternative ways of forming spatial beliefs. You can come to believe that an object is square either by endorsing the content of your experience (in which case your belief is likely to inherit the error in your perception) or by measuring its sides with a ruler.

We like the idea that a Matrix scenario could turn out not to be skeptical at all. But, given our arguments that the real world is a light-weight scenario (at least for many of our spatial experiences), we cannot follow Chalmers and wield the impossibility of light-weight scenarios to deliver the verdict that the matrix is not skeptical. However, a more general spatial functionalism, distinct from the causal-phenomenal variety, retains the potential to deliver the desired result. With the distinctions between skeptical scenarios in hand, we can say more. We are suspicious that things might turn out even worse for us *homo sapiens*. The actual world could turn out to be a middleweight skeptical scenario for spatial experience. The results of Hibbard et al. (2012) suggest that circle-ish experiences are almost never veridical. We are always just a bit off. To establish a full-blooded middleweight scenario, we'd need similar results for all other spatial experiences. But there is certainly cause for concern. The true meta-semantic theory of perceptual representation might rule out heavyweight scenarios, but it will probably not rule out the possibility that we are in a middleweight scenario, and it will certainly not rule out lightweight scenarios.

Despite these worries about actuality, we remain optimistic about the matrix. A matrix could, we suspect, involve more veridicality than the actual world! In the actual world, the normal causes of E_{frontal} and $E_{\text{depth-square}}$ are different. (Reminder: E_{frontal} and $E_{\text{depth-square}}$ are experiences as of squares, presented frontally and in depth respectively. They seem to present the same shape.) But the designers of the matrix might have been kinder to us than Mother Nature has been, setting things up so that E_{frontal} and $E_{\text{depth-square}}$ have the same normal cause. Similarly, they could design the causal structure of our perceptual system so that most of the other perceptual illusions we have discussed do not occur. Such a matrix would be neither a light-, medium-, nor heavy-weight skeptical scenario. This is a surprising and counter-intuitive result. Some of our matrix counterparts do better, with respect to the veridicality of their spatial experiences, than we actually do. These remarks apply, *mutatis mutandis*, to brains in vats and those deceived by Cartesian demons.

9. Conclusion

This paper has examined spatial functionalism, the view that our spatial experiences and spatial concepts pick out whatever property plays a certain functional role. The most prominent version of spatial functionalism alleges that we pick out spatial properties in virtue of their role in causing perceptual experiences (Chalmers 2019). This view faces significant difficulties. Vision science suggests that many of our spatial experiences are subject to lifelong, systematic illusion. Spatial experiences do not always represent the properties that normally

cause them. These arguments do not, however, preclude a version of spatial functionalism that construes the reference-fixing conditions for spatial representations differently.

Chalmers's version of spatial functionalism entails that a certain class of matrix and envatment situations are not in fact skeptical scenarios. Like Chalmers, we are attracted to the idea that some matrix scenarios do not involve systematic deception; matrixes needn't involve substantively more deception than actuality. We disagree about which matrix situations are not skeptical. Furthermore, we conjecture that certain matrix scenarios, with particularly cooperative psychophysics, result in less illusion than the actual world.

While we have not presented a comprehensive rival functionalist proposal that makes these predictions, we have shown that there is quite a lot of space under the spatial functionalist tent. Much of that space is compatible with (a) the twin-earthability of spatial concepts, (b) the possibility and actuality of systematic illusion and lightweight skeptical scenarios, (c) our matrix-counterparts' spatial experiences and concepts representing computer programs or the like, and (d) certain matrix scenarios being neither light-, medium-, nor heavyweight skeptical scenarios.

Vision science places constraints on the meta-semantics of spatial representation. No viable meta-semantics can entail that we are more reliable about spatial properties than we really are. This is where causal phenomenal functionalism founders. In turn, the meta-semantics of spatial representation places constraints on the prospects for antiskepticism. If the true meta-semantics entails that our spatial experiences and concepts could have picked out bits of a computer algorithm, then the matrix need not be a scenario where our spatial representations fail to refer. We think spatial experiences and concepts could have picked out bits of a computer program. But this is not because they pick out whatever property normally causes them. These considerations recommend a type of spatial functionalism that, unlike the causal-phenomenal view, permits lightweight skeptical scenarios but is also compatible with certain matrix and envatment scenarios not being skeptical scenarios of any weight. That is the type of functionalism we hold out for. The view should have a meta-semantics that is properly tamed by the vision science. The view should also deliver an antiskeptical payoff that Chalmers's own view does not, and for which we have argued: There are matrix scenarios where our doppelgängers do better than we do.[29]

[29] We are grateful to audiences at the Pictures, Perception, and Consciousness Workshop at NYU Abu Dhabi and the Philosophy and Psychology of Visual Space Workshop at Ohio State. Thanks also to David Chalmers, to Peter Epstein for detailed comments on an earlier draft of the paper, and to an anonymous referee for this journal.

References

Bayne, T. *The Unity of Consciousness*. Oxford: Oxford University Press, 2010.

Biederman, I. "Recognition-by-Components: A Theory of Human Image Understanding," *Psychological Review* 94(2) (1987): 115–147.

Blumenthal, L. M. *A Modern View of Geometry*. New York, NY: Dover, 1961.

Briscoe, R. E. "Vision, Action, and Make-Perceive," *Mind & Language* 23(4) (2008): 457–497.

Chalmers, D. J. "The Matrix as Metaphysics." *Philosophers Explore the Matrix*, ed. C. Grau. Oxford: Oxford University Press, 2005 (132–176).

Chalmers, D. J. "The Foundations of Two-Dimensional Semantics." *Two-Dimensional Semantics: Foundations and Applications*, eds. Manuel Garcia-Carpintero, Josep Macia. Oxford: Oxford University Press, 2006a (55–140).

Chalmers, D. J. "Two-Dimensional Semantics." *The Oxford Handbook to the Philosophy of Language*, eds. E. Lepore, B. Smith. Oxford: Oxford University Press, 2006b (574–606).

Chalmers, D. J. *Constructing the World*. Oxford: Oxford University Press, 2012.

Chalmers, D. J. *The Character of Consciousness*. New York, Oxford University Press, 2010.

Chalmers, D. J. "Structuralism as a Response to Skepticism," *Journal of Philosophy* 115 (2017): 625–660.

Chalmers, D. J. "Three Puzzles About Spatial Experience." *Blockheads: Essays on Ned Block's Philosophy of Mind and Consciousness*, eds. A. Pautz, D. Stoljar. Cambridge, MA: MIT Press, 2019 (109–138).

Cuijpers, R. H., A. M. L. Kappers, and J. J. Koenderink "Visual Perception of Collinearity," *Perception & Psychophysics* 64(3) (2002): 392–404.

Epstein, P "Shape in a Relativistic Universe," *Mind* 127(506) (2018): 339–379.

Fodor, J. A. *A Theory of Content*. Cambridge, MA: MIT Press, 1990.

Ganson, T. "Sensory Malfunctions, Limitations, and Trade-Offs," *Synthese* 195(4) (2018): 1705–1713.

Green, E. J., and S. Schellenberg "Spatial Perception: The Perspectival Aspect of Perception," *Philosophy Compass* 13(2) (2018): e12472.

Hibbard, P. B. et al. "Misperception of Aspect Ratio in Binocularly Viewed Surfaces," *Vision Research* 70 (2012): 34–43.

Hummel, J. E. "Object recognition." *Oxford Handbook of Cognitive Psychology*, ed. D. Reisburg. Oxford: Oxford University Press, 2013 (32–46).

Hurley, S. L. *Consciousness in Action*. Cambridge, MA: Harvard University Press, 1998.

Lescroart, M. D., and I. Biederman "Cortical Representation of Medial Axis Structure," *Cerebral Cortex* 23 (2013): 629–637.

Lewis, H. R., and C. H. Papadimitriou. *Elements of the Theory of Computation*. Prentice Hall, 1981.

Loomis, J. M. et al. "Visual Space Perception and Visually Directed Action," *Journal of Experimental Psychology: Human Perception and Performance* 18(4) (1992): 906–921.

Loomis, J. M., J. W. Philbeck, and P. Zahorik "Dissociation Between Location and Shape in Visual Space," *Journal of Experimental Psychology: Human Perception and Performance* 28 (5) (2002): 1202–1212.

Masrour, F. "Space Perception, Visual Dissonance and the Fate of Standard Representationalism," *Noûs* 51(3) (2017): 565–593.

McLaughlin, B. P. "The Skewed View from Here: Normal Geometrical Misperception," *Philosophical Topics* 44(2) (2016): 231–299.

Mendelovici, A. "Reliable Misrepresentation and Tracking Theories of Mental Representation," *Philosophical Studies* 165(2) (2013): 421–443.

Millikan, R. G. "Biosemantics," *The Journal of Philosophy* 86(6) (1989): 281–297.

Noë, A. *Action in Perception*. Cambridge, MA: MIT Press, 2004.

Palmer, S. E. *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press, 1999.

Pautz, A. "A Simple View of Consciousness." *The Waning of Materialism*, eds. Robert C. Koons, George Bealer. Oxford: Oxford University Press, 2009 (25–66).

Pizlo, Z. *3D Shape: Its Unique Place in Visual Perception*. Cambridge, MA: MIT Press, 2008.
Rabin, G. O. "How to Twin-Earth a Phenomenal Concept," ms. Available at www.gabrielrabin.com/s/How-To-Twin-Earth.pdf.
Searle, J. *Seeing Things as They Are: A Theory of Perception*. Oxford: Oxford University Press, 2015.
Shapiro, S. *Thinking About Mathematics: The Philosophy of Mathematics*. Oxford: Oxford University Press, 2000.
Stoljar, D. "Chalmers v. Chalmers," *Noûs* (forthcoming).
Thompson, B. "The Spatial Content of Experience," *Philosophy and Phenomenological Research* 81(1) (2010): 146–184.
Tittle, J. S. et al. "Systematic Distortion of Perceived Three-Dimensional Structure from Motion and Binocular Stereopsis," *Journal of Experimental Psychology: Human Perception and Performance* 21(3) (1995): 663–678.
Todd, J. T. "The Visual Perception of 3D Shape," *Trends in Cognitive Sciences* 8(3) (2004): 115–121.
Todd, J. T. et al. "On the Affine Structure of Perceptual Space," *Psychological Science* 12 (3) (2001): 191–196.
Todd, J. T. et al. "The Perception of Doubly Curved Surfaces from Anisotropic Textures," *Psychological Science* 15(1) (2004): 40–46.
Tye, M. "The Problem of Common Sensibles," *Erkenntnis* 66(1–2) (2007): 287–303.
Wagner, M. "The Metric of Visual Space," *Perception & Psychophysics* 38(6) (1985): 483–495.
Wagner, M. *The Geometries of Visual Space*. Mahwah, NJ: Lawrence Erlbaum, 2006.
Wagner, M., and A. J. Gambino "Variations in the Anisotropy and Affine Structure of Visual Space: A Geometry of Visibles with a Third Dimension," *Topoi* 35(2) (2016): 583–598.
Wolfe, U., L. T. Maloney, and M. Tam "Distortions of Perceived Length in the Frontoparallel Plane: Tests of Perspective Theories," *Perception & Psychophysics* 67(6) (2005): 967–979.