Technical Note

Shape, perspective, and what is and is not perceived: Comment on Morales, Bax, and Firestone (2020)

Johannes Burge[1,2,3,*] and Tyler Burge[4,*]

[1] Department of Psychology, University of Pennsylvania

[2] Neuroscience Graduate Group, University of Pennsylvania

[3] Bioengineering Graduate Group, University of Pennsylvania

[4] Department of Philosophy, University of California, Los Angeles

[*] email correspondence on issues in perceptual psychology: jburge@upenn.edu
[*] email correspondence on issues in philosophy of perception: burge@ucla.edu

Abstract

Psychology and philosophy have long reflected on the role of perspective in vision. Since the dawn of modern vision science—roughly, since Helmholtz in the late 1800s—scientific explanations in vision have focused on understanding the computations that transform the sensed retinal image into percepts of the three-dimensional environment. The standard view in the science is that distal properties— viewpoint-independent properties of the environment (object shape) and viewpoint-dependent relational properties (3D orientation relative to the viewer)–are perceptually represented and that properties of the proximal stimulus (in vision, the retinal image) are not. This view is woven into the nature of scientific explanation in perceptual psychology, and has guided impressive advances over the past 150 years. A recently published article suggests that in shape-perception, the standard view must be revised. It argues, on the basis of new empirical data, that a new entity—perspectival shape— should be introduced into scientific explanations of shape perception. Specifically, the article's centrally advertised claim is that, in addition to distal shape, perspectival shape is perceived. We argue that this claim rests on a series of mistakes. Problems in experimental design entail that the article provides no empirical support for any claims regarding either perspective or the perception of shape. There are further problems in scientific reasoning and conceptual development. Detailing these criticisms and explaining how science treats these issues is meant to clarify method and theory, and to improve exchanges between the science and philosophy of perception.

Keywords

Technical Note

The role of perspective in vision has long drawn intense interest in perceptual psychology, philosophy, and art. Currently, nearly everyone in these disciplines holds that we perceive environmental objects and their properties, and that we do not perceive the images that they cast on the retina. However, we are also consciously aware of certain properties that depend on perspective (or viewpoint). For example, when accurately perceiving a rotated circular dinner plate as circular, one has a shape-awareness similar to the shape-awareness one has while accurately perceiving a head-on elliptical plate as elliptical. It may be tempting to assert that we perceive all that we are consciously aware of. It may also be tempting to assert that there is no meaningful distinction between what is perceived and what is merely sensed. But there are scientific reasons to reject both assertions. This conscious shape-awareness—often called 'painter's perspective' (Alberti, 1435; Kemp, 1990)—is related to but is not the same as perceiving. The great Renaissance painters mastered this perspective, with some difficulty, and used it to produce realist works of art. Painter's perspective has led to a long-running discussion in the history of philosophy. Many scientists are not clear, off-hand, about how to describe the situation. But the outline of a correct position on shape-perception, and perspective, has been present in vision science for decades.

A recent article—Morales, Bax, and Firestone (2020), 'Sustained representation of perspectival shape', PNAS, 117(26), 14873-14882—reports a series of experiments that purport to bear on these matters. The authors' *centrally advertised claim* is that, in addition to distal shapes of objects in the environment, 'perspectival shapes are represented by mechanisms of perception'. A natural interpretation takes this claim to imply that humans see (accurately represent in visual perception) perspectival shape. The authors write that either visual 'perception of [a rotated dinner plate] switches back and forth between its distal and perspectival shapes… almost like a bistable figure' or there are 'coexisting distal and perspectival interpretations'. 'Perspectival shape' is not a standard term in vision science. The authors never explain their intended meaning for the term. It is clear that perspectival shape is determined by the distal shape and its 3D orientation relative to the viewer (or equivalently, by the distal-shape silhouette from the viewer's viewpoint). But it is unclear whether the authors intend the term to refer to a property of the light-field located between the distal shape and the retinal image, the projected shape in the retinal image (*retinally projected shape*), the retinal or cortical encodings of retinally projected shape, or some other shape property. For many points that immediately follow, it will not matter what exactly the authors intend by the term. (We later revisit this issue.) In any case, the claim that perspectival shape is perceived, or is something 'interpreted' in perception, departs from scientific consensus. So it is important to examine the basis for the claim.

Scientific consensus is as follows: Distal properties are perceptually represented. Distal properties include viewpoint-independent properties, like object shape and object size, and viewpoint-dependent relational properties, like 3D orientation and distance relative to the viewer. The retinal image and features in it—including retinally projected shape and size—are not perceived. The retinal image and its features do, however, affect and retain a presence in—that is, cause viewpoint-dependent analogs in—perceptual states (see *Vision Science* below). Perspectival shape, *construed as distinct from retinally projected shape and its effects*, does not figure in scientific explanations of perception, or in scientific accounts of perceptual performance.

Morales et al. tasked subjects with rapidly discriminating distal circles from distal ellipses in two conditions. In the easy condition, the distal circle and distal ellipse were both head-on (not rotated in depth), resulting in projected shapes in the retinal image that were substantially different (i.e. circular vs. elliptical). In the hard condition, the distal circle was rotated so that its projected shape in the retinal image matched the elliptical projected shape of the head-on ellipse (Fig. 1A). In each of nine experiments—all of which shared a similar design—human subjects responded more slowly in the hard condition than in the easy condition, and performed with very high overall accuracy ($\geq$95%). The authors write that their speed and accuracy results 'cannot be explained by low-level properties [of the retinal image]', are explained by 'perspectival shape per se (rather than some other factor)', are 'inconsistent with claims that we perceive distal properties only', and demand a 'considerably expanded role for perspectival shape than is traditionally favored'.

The authors' results provide no empirical support at all for the above-quoted assertions, or the centrally advertised claim. The article's many shortcomings are rooted in poorly controlled and poorly conceived experiments, failure to report in the article relevant data from their experiments, and faulty reasoning used to evaluate the data that was reported. These flaws prevent the authors from drawing *any* justified conclusions from their data regarding perspective or shape-perception. Additional points of concern fall into three main categories. First, the authors never explain what role the retinal image plays in their account. They never explain how to experimentally distinguish effects of perspectival shape from those of the retinal image, or how perspectival shape would contribute to explanation beyond the already recognized contribution of the retinal image. As noted, they do not even explain what perspectival shape is. Second, the article uses imprecise concepts, unclearly formulated hypotheses, shifting—often non-standard—terminology, and misleading contrasts with other positions. Third, the authors substantially misrepresent mainstream views in vision science and do not discuss relevant, scientifically informed philosophy of perception. These misrepresentations make their experiments look more innovative than they are, and their central claim appear more appealing than it is.

Experimental Design and Results
Experiment 1 presented graphically-rendered distal circles and ellipses to subjects and established the basic result. Namely, that head-on ellipses and rotated circles that generate projected shapes in the retinal image with identical aspect ratios take longer to discriminate than circles and ellipses that generate projected shapes with different aspect ratios. The authors assume that the observed response slowdowns are due to *interference*, writing that their results 'provide… evidence that the interference is caused by the representational similarity produced by similar perspectival shapes per se'. The authors take Experiment 1 to provide initial evidence for their advertised claim. Aware that subjects could use some other low-level 'visual feature, such as area or width' in the retinal image as a response-heuristic, the authors attempted to design subsequent experiments to 'rule out this and other similar strategies'. The structure of the article and the authors' own statements make clear that they believe support of their claim rests on success in doing so. The authors do not succeed. They do not rule out the possibility that low-level retinal-image features other than shape account for the data. Indeed, using their methods, they *cannot* control for all low-level retinal image features other than shape, while retaining above-chance distal shape discrimination performance in the hard condition. We explain below.

Experiment 2 manipulated the area and width of the projected shapes in the retinal images while keeping aspect ratio constant. The original findings were replicated. The authors write that these new results 'suggested that it was the similarity or dissimilarity in perspectival shape per se (rather than some other factor) that drives these response-time differences'. After additional experiments, they conclude that the results '…cannot be explained by low-level properties [of the retinal image]…'. This conclusion is unjustified. Uncontrolled low-level image properties were present in all experiments (Fig. 1B). For example, in the hard conditions, subjects could identify the distal circle as the stimulus with the darker shading, or with the visible edge, among other cues (see Fig. 1AB). It is not true, as the authors write, that theirs is a '[paradigm] where targets and distractors differ along just one feature (as ours do in only shape)'. The incomplete controls prevent the authors from justifiably attributing response-slowdowns to perspectival (or retinal) shape similarities. The article provides no evidentiary basis for invoking anything but the uncontrolled retinal image cues in explanation of the results. Shape similarities could, in principle, be among the factors. But the authors' experiments have not shown that they are.
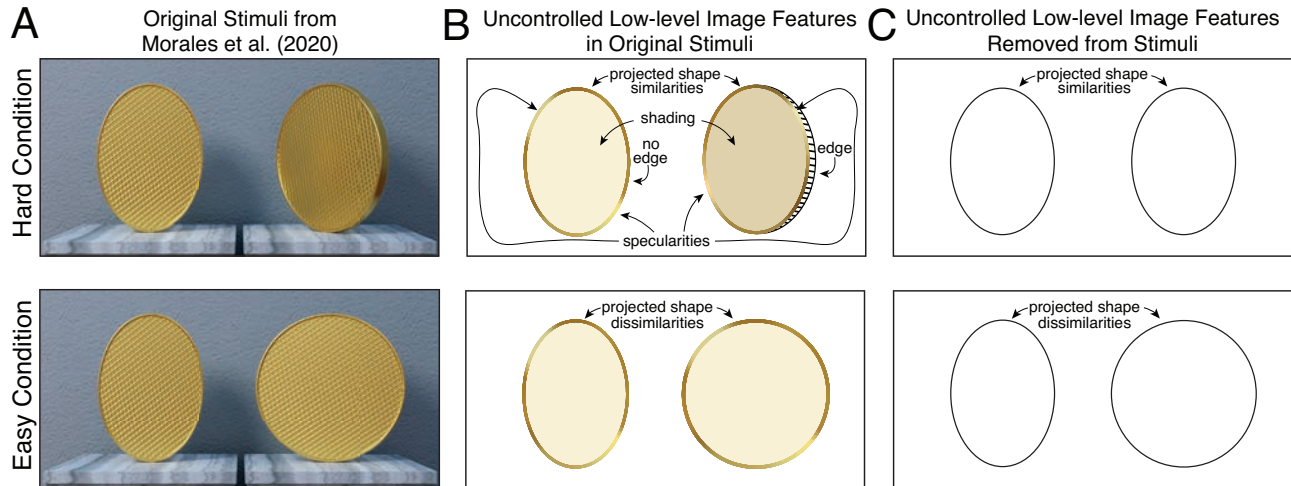
Figure 1. Experimental stimuli from Morales et al. (2020) and problems with experimental design. **A** Original stimuli. In the hard condition (top), subjects were tasked with discriminating a distal head-on ellipse (left) from a distal rotated circle (right) that gave rise to images with identical aspect ratios. In the easy condition (bottom), subjects were tasked with discriminating a head-on distal ellipse (left) from a head-on distal circle (right). Distal shape discrimination and identification performance was good in both conditions. But subjects were slower, and less accurate, in the hard than in the easy condition. The authors assert that 'similarity… in perspectival shape per se (rather than some other factor)' is the cause of response slowdowns in their task. **B** Uncontrolled low-level image features in the original stimuli. In the easy condition, subjects could discriminate distal shape on the basis of the projected shape in the retinal image (elliptical vs. circular). In the hard condition, because projected shapes were matched, subjects had to discriminate distal shapes based on differences in shading, specularities, and edge visibility, among other cues: weaker cues to distal-shape differences in this context. The fact that these and other cues (see main text) differ between the hard and easy conditions entails that response slowdowns and accuracy decreases in the hard condition cannot be justifiably attributed to projected shape similarities. Projected shape similarities may be involved, but the experiments do not show that they are. **C** Controlling all low-level image features renders discrimination of distal shape impossible in the hard condition. Experiments isolating projected shape similarities and dissimilarities as the only cues differentiating the easy and hard conditions could not have yielded above-chance distal shape discrimination in the hard condition. (Images in A adapted from Morales et al., 2020.)

Psychophysics attributes performance differences to a factor only if that factor has been isolated as the cause of the performance differences in the relevant conditions. Isolating shape similarities—perspectival or retinal—as causing the response-time differences would have required that all other low-level features (cues to distal shape in the retinal image) be matched in the hard conditions. If all low-level image features were matched (two elliptical outlines with the same aspect ratio that lack—or otherwise do not differ in—shading, specularities, edge visibility, texture, binocular disparity, focus-cues, motion parallax, etc.), cues to distal-shape differences would be eliminated (Fig. 1C). Under such circumstances, the two distal shapes would be impossible to discriminate. The authors could not have succeeded at the task they set themselves—to rule out all low-level retinal image features other than shape. The experiments were not only poorly controlled. They were uncontrollable.

These experimental design-failures undermine all the authors' claims about perspectival shape (see Fig. 1AB). Because of unsuccessful design, the experimental results provide no grounds upon which to conclude that interference by perspectival (or retinal) shape similarities—or by 'representational similarit[ies]' of *any kind*—are responsible for reaction-time slowdowns. The experiments do not demonstrate interference by anything at all.

The retinal-image cues available for discriminating distal ellipses and circles, including retinally projected shape, are as a group weaker and less reliable in the hard condition than in the easy condition. Distal circles and ellipses that generate retinally projected shapes with matched aspect ratios force subjects to rely more on other cues (e.g. shading, specularities, edge visibility; see Fig. 1AB). Less informative, less reliable cues tend to cause slower response times and less accurate discriminations.

The authors several times write as if their experimental results are out of step with common findings regarding the relationship between speed and accuracy. But there is nothing surprising about their results. The response-time slowdowns in the hard conditions are accompanied by accuracy degradations, results not reported in the article (Fig. 2AB). In Experiment 1, the authors report 97% accurate discriminations of distal shape. They do not report that discriminations in the hard condition were less accurate than those in the easy condition. This was true both for the filtered data analyzed in the article (94.6% vs. 98.5%) and for the raw data (83.5% vs. 87.1%). Experiments 5 and 6 added dynamic cues and delayed responding in order to, according to the authors, 'emphasize the 3D shape of the object… thereby [making] any influence of perspectival shape [on response times] that much more surprising'. The authors point to these results as particularly compelling. But these results also exhibit accuracy differences (filtered data: 91.9% vs. 98.2% and 93.7% vs. 96.3%). Experiments 8 and 9 used real, non-graphically rendered stimuli. The authors write 'if perspectival shapes interfere here too, this would provide extremely strong and ecologically valid evidence that perspectival shapes are represented by mechanisms of perception'. But these experiments introduced more uncontrolled cues (binocular disparity & motion parallax). And, here too, there are accuracy differences between the hard and easy conditions (filtered data: 94.5% vs. 97.4% and 97.2% vs. 98.4%). All nine experiments follow the same pattern (Fig. 2B).
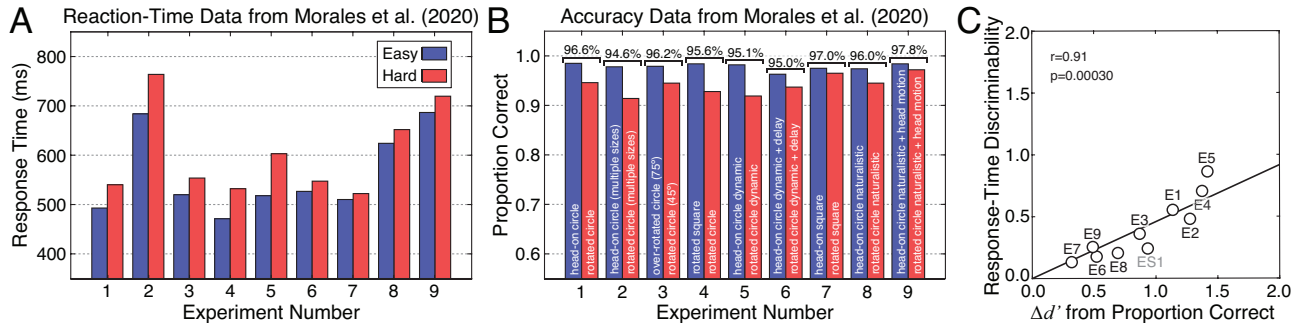
Figure 2. Reaction-time and accuracy data from Morales et al. (2020). The original article reported filtered reaction-time data separately for the easy and hard conditions, and accuracy data averaged across—but not separately for—the easy and hard conditions in nine experiments. **A** Reaction-time data reported in the original article. In each experiment, subjects were slower to identify the head-on ellipse when paired with a rotated (red) vs. a head-on (blue) circle (see Fig. 1A, top vs. bottom). Despite numerous uncontrolled image cues (see Fig. 1AB) the authors conclude: 'Across these many studies [Exps. 1-7], only the representational similarity [explicitly taken to be perspectival shape-similarity] between the [head-on elliptical] target and the rotated [circular] distractor explained the observed RT slowdown'. **B** Accuracy data in the easy and hard conditions. Subjects were always less accurate in the hard (red) than in the easy (blue) condition. Morales et al. (2020) reported average accuracy (brackets above bars), but did not report accuracy separately for the easy and hard conditions. It is not the case, as the authors write, that because 'accuracy [averaged across easy and hard conditions]… was near ceiling… that subjects were not confused at all about the true shapes of the stimuli'. **C** A strong relationship exists between reaction-time data and accuracy data across the experiments. There is a significant correlation (r=0.91; p=3.0x10⁻⁴) between the discriminability of response-time data in the easy and hard conditions, on one hand, and differences in sensitivity (d-prime) calculated from the accuracy data in the two conditions, on the other (see Methods). The 95% bootstrapped confidence interval ranges from 0.81 to 0.98. The labels (E1, E2, etc.) indicate the experiment corresponding to each data point. The gray data point (ES1) is from a supplemental experiment (Morales, Bax, & Firestone, 2021), performed in response to another published criticism (Linton, 2021).

> The following statements by the authors are seriously misleading regarding their own data:
>
> > … accuracy on the primary task was near ceiling. This strongly suggests that subjects were not confused at all about the true shapes of the stimuli…
>
> > Perspectival shapes are sufficiently present in perception [to slow response times while] at the same time, [perspectival shapes] are sufficiently distinct from the distal contents of our shape representations that they do not confuse subjects about the true properties of objects.

The unreported accuracy data shows that subjects were less accurate and, in this sense, more often confused in the hard conditions.[1]

We examined whether there was a quantitative relationship between the reaction-time and accuracy data. The discriminability of the reaction-time data between the easy and hard conditions is tightly correlated with sensitivity differences computed from the accuracy data (i.e. d-prime differences between the easy and hard conditions; see Methods, Fig. 2C). The close relationship between speed

---

[1] We emphasize that even if there were no relationship between the speed and accuracy results (which *would* be a bit surprising), the experiments as designed would still have provided no basis for thinking that interference from shape similarities of any kind caused response-time slowdowns in the hard conditions, because of the uncontrolled cues.

and accuracy has been known for decades, and is predicted by the drift-diffusion model, a model commonly used to account for such data (Barack & Gold, 2016; Norman & Wickelgren, 1969; Palmer, Huk, & Shadlen, 2005; Uchida, Kepecs, & Mainen, 2006; Watson, 1979; Wickelgren, 1977). The reported response-time and accuracy differences accord with well-known phenomena. The disassociation between speed and accuracy implied by the quoted text does not exist. The authors' advertised claim—which purports to explain the supposed disassociation—has nothing to explain.

Scientific Reasoning

Two main criticisms have been advanced so far. First, because of poor experimental design, Morales et al. (2020) can support no claims about perspective or about shape-representation in vision. Second, the tight relationship between the speed and accuracy results is in keeping with a huge body of experimental work, and does not call for rethinking standard explanations. The remainder of our article focuses not on experimental issues, but on conceptual, theoretical, and terminological issues.

Morales et al. (2020) is unspecific about what perspectival shape is. We discuss how, for each of several conceptions of perspectival shape that appear consistent with different passages in the text, the scientific reasoning is problematic.

One conception locates perspectival shape in some intermediate position(s) between the distal shape and the retinal image. Such an intermediate entity could be a collection of 3D directions, or some other property of the light-field, that is jointly determined by distal shape and its 3D orientation relative to viewer. Such an intermediate entity, whatever the specifics, must form the same silhouette from the viewer's viewpoint that the distal shape does, and must result in the same retinally projected shape. Some contemporary philosophers—influenced by Locke (1690)—hold that this sort of intermediate entity is perceived (see *Philosophy* below). Vision scientists will tend to find this view odd.

We give precedence to this conception because the authors once clearly distinguish perspectival shape from the retinal image, because perspectival shape is not distal shape, and because it seems a plausible reading of what the authors believe. We show that, on this conception, their advertised claim invokes an entity that has played no role in vision science on the basis of data that provides no evidence to support that invocation.

Science gives a role to a new potential explanatory factor only when the new factor is shown to play an explanatory role not played by factors already cited in scientific explanation. To test whether perspectival shape, conceived as an intermediate entity, is perceived or has any other role in scientific explanation, experiments would have to disentangle its role from the role of retinally projected shape. In humans, such experimental disentanglement is likely impossible. Many perspectival and retinal-image properties, including shape, are perfectly coupled in the human visual system because the

distance between the lens and the retina is fixed.[2] Hence, experiments with humans could not show that perspectival shape does explanatory work over and above work already done by the retinal image. Because of the well-established role of the retinal image and the disentanglement problem just described, the retinal image obviates any explanatory role for perspectival shape. So the claim that perspectival shape is perceived, or plays any other explanatory role, adds nothing to scientific explanation. Vision science does not, and does not need to, refer to intermediate entities. Morales et al. do not propose an explanatory role for perspectival shape that is distinct from the role already played by projected shape in the retinal image. There is no reason to give it a role. Hence, in the spirit of Ockham's razor, invocation of perspectival shape should be rejected (see also Lande, 2018).

The authors never substantively discuss the central source of viewpoint-dependent visual information: the retinal image. They once mention, in a footnote, that their results do not adjudicate between interference by perspectival shape and interference by the retinal image. (Recall that the results do not constitute evidence for interference by anything.) They write,

> … the results here indicate such representational similarity without specifying the dimension of such similarity, or the specific features that ground this similarity... it may be important to distinguish between interference caused by matching perspectival shapes vs. by persisting retinal images... Our results here cannot adjudicate between these extremely subtle options...[3]

The option between invoking perspectival shape, conceived as an intermediate entity, and invoking the retinal image to account for their results is not 'extremely subtle'. The option lies at the heart of the difference between the authors' advertised role for perspectival shape and a standard explanation that they count incompatible but present no evidence against. Independently of problems already discussed, the authors' admission that their results cannot distinguish between interference by perspectival shape and interference by the retinal image, in itself, undercuts the advertisement of 'a

---

[2] Disentanglement of perspectival size, another perspectival property, may be possible in other organisms. In certain fish, the distance between the lens and the retina changes actively (Land & Nilsson, 2002). Because of these changes, retinal-image size (millimeters on the retina) can be decoupled from perspectival size (angular subtense). By presenting stimuli contingent on the distance between the lens and retina, perspectival and retinal-image size could perhaps be experimentally disentangled. We know of no data that supports taking perspectival size or shape to have a role in scientific explanation even for these organisms. But at least relevant hypotheses could be tested.

[3] Repeatedly, throughout the main text, the authors attribute the speed and accuracy results to 'representational similarit[ies] produced by similar perspectival shapes per se', and advertise a new role for perspectival shape (note the title of the original article). Here, they attribute the results to '[unspecified] representational similarit[ies]'. In a reply to Linton (2021), they call this their most 'cautious' conclusion (Morales et al., 2021). The experiments no more support the conclusion that unspecified representational similarities account for the results, than they support the conclusion that perspectival shape similarities do.

considerably expanded role for perspectival shape than is traditionally favored'.[4] This fact is obscured by their unclear formulations, by their implying, mistakenly, that mainstream vision science cannot account for their speed and accuracy results, and by their portraying vision science as maintaining that 'distal [shape is seen] in ways that completely separate [it] from our point of view' (a mischaracterization that we discuss below).

Consider a possible objection: The difference between the standard view and the authors' advertisement that perspectival shape is perceived just amounts to semantics; the difference hinges merely on the "definition" of what counts as perceiving something. Does the difference come to anything more than different stipulations about what things are perceived? It does. Definitions in science are more often substantive than stipulative. They usually come at the end of scientific inquiry, not at the beginning (T. Burge, 2014a; Putnam, 1962; 1973; 1975). They function to capture basic knowledge that derives from scientific investigation. Science's account of what is—and is not—perceived is not stipulative. It is based on evidence. The retinal image initiates a causal process of sensory-state-formation that culminates in percepts of the distal environment. There is no serious debate over whether the retinal image is perceived. The retinal image does, however, contribute to the viewpoint-dependent nature of percepts. Shape percepts that represent a given distal shape with different 3D orientations relative to the viewer are different percepts. Such percepts are commonly associated with different processing speeds and degrees of accuracy (see Fig. 2C). Because (i) differences between these percepts partly consist in analogs of—or, better, sometimes biased encodings of—different retinally projected shapes, because (ii) the effect on perception of perspectival shape does not differ from that of retinally projected shape, and because (iii) we do not perceive the retinal image (or projected shape in it), the claim that we perceive perspectival shape has no empirical basis. As indicated, scientific practice avoids invoking a new entity in scientific theory, if that entity does no new explanatory work. Here, invoking perspectival shape, conceived as an intermediate entity, contributes nothing beyond what is already contributed by invoking the retinal image in explaining how percepts are formed and what is perceived. So the claim that perspectival shape is perceived should be rejected.

Earlier we stated that, however 'perspectival shape' is construed, the claim that perspectival shape is perceived has no merit. We have considered perspectival shape as an intermediate entity. Other possibilities are retinally projected shape, its subsequent encodings in the retina or the brain, and the union of any of the discussed construals.

---

[4] The authors' acknowledgment that they cannot distinguish perspectival shape from the retinal image as the cause of their results, in effect, concedes that there may be nothing more going on than the well-appreciated fact that properties of the retinal image (perhaps but not necessarily including retinally projected shape) can impact the speed and accuracy of percept-formation.

Suppose that the authors intend 'perspectival shape' to refer to retinally projected shape. (We think this construal unlikely because they distinguish the two in the footnote quoted above.) Then their advertised claim implies that the retinal image is perceived. Such a view conflicts with basic tenets of the science, is subject to an obvious variant of the Ockham-like argument above, and is also supported by none of their data. (The same argument holds for taking perspectival shape to be, or to include, encodings of the retinal image.) Suppose instead that, in writing '… perspectival shapes not only exist as images projected on our eyes but also guide mechanisms of perception…', they intend to claim that retinally projected shapes (and/or subsequent encodings of them) are sensed, but not perceived. Then their claim simply rebrands the well-appreciated fact that the retinal image impacts the speed and accuracy of percept-formation. So, on all these construals (or the union of any of them), the authors have not demonstrated a need for an 'expanded role for perspectival shape'. Perspectival shape—construed as an intermediate entity–is not absent from vision science because of neglect. It is absent because it is unnecessary for scientific explanation.

We stress that all of the points made in this section are independent of the fact that poor experimental controls prevent justifiable attribution of the speed and accuracy results to perspectival shape similarities—or to 'representational similarit[ies]'—of any kind.

Conceptual and Terminological Distinctions

Lying behind vision science's treatment of perspective is a distinction between sensing and perceiving. The science takes as basic that distal properties like distal shape are perceived. It also takes as basic that properties of the retinal image like retinally projected shape are not perceived. They are, of course, sensed. Explanations in mainstream vision science ultimately center on a computational sequence of states running from the sensed retinal image to perceptual representations of distal properties (Brainard & Freeman, 1997; J. Burge, 2020; J. Burge & Geisler, 2011; 2014; 2015; Chin & Burge, 2020; Cottaris, Wandell, Rieke, & Brainard, 2020; Erdogan & Jacobs, 2017; Geisler, 1989; Jaini & Burge, 2017; Kersten, Mamassian, & Yuille, 2004; Knill & Richards, 1996; Kupers, Carrasco, & Winawer, 2019; Kwon, Tadin, & Knill, 2015; Marr, 1982; Murray, Boyaci, & Kersten, 2006; Storrs, Anderson, & Fleming, 2021). These computations underlie perceptual constancies, capacities to stably represent distal properties given very different retinal images. The fact that scientific explanation relies on the distinction between perceiving, on one hand, and sensing that is not perceiving (*mere sensing*), on the other, is widely acknowledged in vision science. But the distinction is often not marked clearly by terminology. This lack of terminological clarity can cloud thinking and make productive, efficient communication within each discipline—and across disciplines—more difficult. The points that follow are intended primarily to encourage more rigorous, standardized use of terminology.

The article's imprecise, inconstant use of language generates ambiguous, unclear theoretical positions. For example, no distinction is made among (i) 'represents', (ii) 'is sensitive to', 'carries information about', 'reflects', 'tracks', is 'impacted by', and (iii) 'perceptually represents', 'perceives', 'sees'. The unqualified term 'represents' is ambiguous as between (ii), terms associated with mere sensing, and (iii), terms associated with perception. So mere sensing is not clearly distinguished from perception.

Among many unspecific formulations, the authors once write that the 'mind is sensitive to' information deriving from the projective similarity between the rotated circular coin and the head-on elliptical coin. The discussion slides from this correct point, which is not new, to the advertised claim, which is not correct. The running together of (i)-(iii) also masks the difference between the advertised claim—that we perceive two shapes, perspectival and distal—and the standard view that we perceptually represent distal shape and viewpoint-dependent relations to it like 3D orientation, and are sensitive to (but do not perceive) the shape of the projection in the retinal image.

In many passages, the authors describe their primary issue in vague, unsound ways. They write,

> On one hand, some philosophers sympathize with the dominant view in psychology and vision science, arguing that perceptual experiences are primarily or exclusively about distal environmental properties. At the same time, however, many other philosophers dissent, arguing that our visual experiences are better described by a "dual" character, such that perceptual experience reflects both the true distal properties of objects and their perspectival properties–a circle and an ellipse at the "same time" as it were.

The passage takes two compatible views to be incompatible. There is nothing incompatible between 'perceptual experiences [that] are primarily... about distal environmental properties' and perceptual experiences that 'reflect... both distal properties of objects and their perspectival properties'.

Further, in context, the passage seems to imply that because a viewpoint-dependent shape property (perspectival shape or retinally projected shape) is 'reflect[ed]' in 'perceptual experience', that it is perceived. The term 'perceptual experience' tends to suggest conscious awareness. Conscious awareness and perception often co-occur, but they are not the same (T. Burge, 2014a). Conscious awareness of a property in perceptual experience does not entail perception of the property. Percepts of things can, in principle, be conscious or unconscious. Mere sensings of things can be unconscious or conscious. Introspecting, or otherwise determining, what one is consciously aware of cannot distinguish what is perceived from what is merely sensed. Science discovers what we perceive–what we perceptually estimate or perceptually categorize—by giving causal explanations, not by engaging

in introspection. The scientific discovery process has established different roles for perceiving and mere sensing.[5]

Science does not yet explain why or predict which particular aspects of sensory or perceptual states are conscious. These are tasks for an extension of current science (Denison, Block, & Samaha, 2021).

Representing something—perceptually or otherwise—always involves representing it in a specific way. What is represented (here, distal shape and its 3D orientation relative to the viewer) is distinct from how it is represented (here, in a specific viewpoint-dependent way). (The distinction between what and how something is represented is well-known in philosophy.) The specific viewpoint-dependent way in which percepts represent things—the *how* aspect of representation—is controlled by the retinal image and includes analogs of retinal-image features like retinally projected shape (T. Burge, 2022). But no retinally projected property is perceived, is perceptually represented, or perceptually represents (or estimates) anything on its own.

Vision Science
The article's characterizations of vision science are inaccurate. The article represents the orthodox view in psychology and neuroscience of perception to be that we 'see the distal properties of objects in ways that completely separate them from our point of view'. It portrays mainstream vision science as holding that viewpoint-dependent information fails to 'persist in', is 'separate[d]' from, is not 'tracked' in perception, and is 'truly absent from the object representations that inform perception and action'. Not one of these views is orthodox in vision science. The article's characterizations are accompanied by passages quoted from the work of vision scientists. We checked with some of the quoted scientists. In each case, they hold views very different from the present authors' portrayals. Vision science has no serious division over whether viewpoint-dependent features of the retinal-image cause viewpoint-dependent analogs to be present in perception. In this sense, retinal-image features '"sustain" a presence in perception.

The entire field of visual psychophysics is based on the fact that differences in retinal image properties, and sensitivity to those differences, impact later visual processing, perception, and perception-driven responses. A huge amount of work has examined perceptual representation (and neural correlates of perceptual representation) of distal, viewpoint-dependent properties like 3D orientation and distance relative to the viewer (Adams et al., 2016; Backus, Banks, van Ee, & Crowell,

---

[5] We emphasize that a theoretical explication of the distinction between mere sensing and perception is not necessary to make the points just made. But, just to provide rough orientation, perception is distinguished by its estimating and categorizing things, and is marked by perceptual constancies, whereas, in vision, mere sensing is centered on encoding information in the retinal image (T. Burge, 2010; 2014a; 2022).

1999; J. Burge, McCann, & Geisler, 2016; Fleming, Holtmann-Rice, & Bulthoff, 2011; Gillam, 1967; Hillis, Watt, Landy, & Banks, 2004; Kim & Burge, 2018; 2020; Knill, 1998; Kuroda, 1971; Malik & Rosenholtz, 1997; McCann, Hayhoe, & Geisler, 2018; Mooney, Marlow, & Anderson, 2019; Ogle, 1950; Rosenberg, Cowan, & Angelaki, 2013; Stevens, 1983; Todd, 2004; von Helmholtz, 1909; Welchman, Deubelius, Conrad, Bülthoff, & Kourtzi, 2005). These properties are not 'completely separate… from our point of view'. Vision scientists routinely recognize that information about point of view suffuses perception. Vision-initiated grasping of a dinner plate, for example, obviously requires that the percept be sensitive to retinal-image cues that enable perceptual representation of the plate's shape, distance, and 3D orientation (Knill & Kersten, 2004; Mamassian, 1997; Vishwanath, Girshick, & Banks, 2005).

To support their characterization of vision research, the authors quote the following passage from Murray et al. (2006) regarding the flow of neural activity as visual processing proceeds through the hierarchy of cortical areas: "the visual system progressively transforms information from a retinal to an object-centered reference frame whereby retinal (information) is progressively removed from the representation". In higher cortical areas, neural activity does tend to correlate more explicitly with the distal stimulus than with the proximal stimulus. Distal properties become progressively easier to decode (for example, using linear methods) from neural activity in higher than in lower visual areas. But viewpoint-dependent information is absent neither from higher cortical areas (Albright & Desimone, 1987; Arcaro, McMains, Singer, & Kastner, 2009; DiCarlo & Maunsell, 2003; Kay, Weiner, & Grill-Spector, 2015; Knapen, 2021; Larsson & Heeger, 2006; Le, Witthoft, Ben-Shachar, & Wandell, 2017; Mackey, Winawer, & Curtis, 2017; Silver & Kastner, 2009; Wandell & Winawer, 2011), nor from perception. Morales et al. over-interpret the point that certain things are 'progressively removed' in higher visual areas to mean that all viewpoint-dependent information is 'absent' or 'discarded' from percepts. More centrally, they conflate neural representation with perceptual representation. That is, they conflate neural activity that correlates well with something, on one hand, with perceptual representation of something, on the other. This conflation is perhaps invited by some of Murray et al.'s terminology, but the distinction is clear to the primary readership of the article. Clarifying the many usages of the term 'representation', and being thoughtful about whether a given usage applies in a given context, is an important topic in need of attention in neuroscience and related fields (Baker, Lansdell, & Kording, 2021; Barack & Krakauer, 2021; T. Burge, 2010; 2014a, 2022).

The authors' mischaracterization of vision science as holding that we 'see the distal properties of objects in ways that completely separate them from our point of view' is used to motivate the suggestion that vision science does not recognize that perception has a "dual character". But perceptual states are widely appreciated to have a dual character in several respects. Not one of these respects separates perception from point of view. In each duality, a viewpoint-independent factor is associated

with a viewpoint-dependent factor. First, each perceptual state represents both distal viewpoint-independent properties (e.g. object-shape) and distal viewpoint-dependent properties (e.g. 3D orientation relative to point of view). Second, each perceptual state represents *what* it represents (re-read the previous sentence) in a specific retinal-image-controlled, viewpoint-dependent *way* (see *Conceptual and Terminological Distinctions* above). Third, perceptual states often support two shape-awarenesses. A rotated distal circular shape appears circular in one sense, while it appears elliptical in another. The elliptical appearance corresponds to the elliptical retinal projection of the rotated circular shape. Few vision scientists would deny that percepts or associated appearances are marked by point of view in any of these three respects.

The authors, of course, go beyond misdescribing vision science's views on perception's dual character. In places, they take appearance similarities of the sort just described to ground explanation of their experimental results. They assert,

> The explanation of [the speed and accuracy] results seems clear and straightforward: [Head-on ellipses are] harder to distinguish from [rotated circles] because the two objects *appear to have something in common*… More precisely, it can be said that they bear a representational similarity to one another.[6]

We agree, and most vision scientists would agree, that in this circumstance, appearance similarities exist. But the authors have not shown that such similarities play any role in accounting for their experimental results. And they move without argument or explanation from there being an appearance similarity to there being a representational similarity—and from there being a representational similarity to there being a similarity in what is represented.

The authors even take dual character to consist of either simultaneous 'coexisting' or alternating 'bistable' *perception of* two shapes: one distal, the other perspectival. This interpretation of dual character really *is* incompatible with views in vision science. It is also unsupported by evidence.

---

[6] Perhaps, in the hard condition, the contributions of similar retinally projected shapes to *how* distal shapes are perceptually represented (see *Vision Science*) could count as 'representational similarit[ies]'. But if this were all that the authors intended, the claim would be obvious. That is, if the authors were to claim merely that there is a viewpoint-dependent shape-similarity in how distal shapes are perceptually represented (or in how they appear), the authors would be stating something that is already well-appreciated (see also Morales et al., 2021).

  If they are claiming something more, as they do throughout much of the article,—namely, that during perception of a rotated distal circle, the elliptical quality that we are aware of is something that is perceived—, then their claim is flatly incorrect. It is either tantamount to the claim that the retinal image, features in it, and/or encodings of it are perceived, or to the claim that some entity intermediate between the retinal image and the distal shape is perceived. The specifics of the claim depend on how 'perspectival shape' is construed (see *Scientific Reasoning*).

  In any case, the experiments give no reason to believe that similarities in how distal shapes are represented (or in how they appear) are involved in explaining the speed and accuracy results. So the obvious claim provides no better explanation of the results than does the incorrect claim that perspectival shape is perceived (see Footnote 3 above).

Philosophy

The authors' descriptions of issues in philosophy lack perspective. Some 17th-18th century philosophers did think that only intermediate perspectival properties (Locke, 1690) or only internal counterparts of such properties (Hume, 1748) are perceived. Remarkably, both Hume and Locke thought that no distal properties are perceived. (Recall that we explicate 'distal properties' in the *Abstract*.) The Humean position that only internal states are perceived is now universally rejected. Although nearly everyone also now rejects the Hume-Locke position that no distal properties are perceived, some contemporary philosophers hold, like Locke, that we perceive intermediate perspectival properties. Philosophical defenses of this view are not scientifically informed. The arguments do not take scientifically acceptable forms (J. Cohen, 2010; Hill, 2016; Schellenberg, 2008). A scientifically acceptable argument for holding that humans perceptually represent perspectival shape, conceived as an intermediate entity (see *Scientific Reasoning*), would have to take a specific form. It would have to describe current scientific explanations of shape-perception, and recognize that they make no reference to perspectival shape. It would have to cite evidence that those explanations cannot account for. And it would have to show that the evidence supports appeal to perspectival shape. Philosophical defenses of the neo-Lockean view do not make such an argument. Typically, they do not even correctly describe how science explains shape-perception. Given that the retinal image does all the work that perspectival shape could do in grounding explanation of perspective in human shape-perception, we think that no scientifically acceptable philosophical defense is possible. Morales et al. barely discuss philosophy that utilizes science's treatment of perspective in perception. It is not correct that 'contemporary [philosophy]' on 'whether… perspectival properties are ever truly "perceived"… only appeals to introspection'. The idea that the role of perspective in visual perception had been, before the authors' work, 'only amenable to argument based on introspection' is simply not true (T. Burge, 2010; 2014b; 2014c; Lande, 2018; Rescorla, 2014).

Throughout their article, Morales et al. shift between using the term 'perception' and using the terms 'perceptual experience' or 'conscious experience'. Vacillation on these matters is also common in philosophy. It would be well to differentiate these terms and topics more rigorously. Morales et al. sometimes seem to be trying *both* to provide insight into conscious awareness of viewpoint-dependent shape properties, and to argue that perspectival shape should have a distinctive role in explanation of percept formation. But there is nothing in their results that provides scientific insight into the nature of conscious awareness, or its role in distal shape estimation and discrimination. Because their experiments were not—and could not have been—successfully controlled, their results show no more about conscious awareness, or conscious experience during perception, than they show about percept-formation. Further, response times were too slow (500-800ms) to discriminate when or where, in

processing, conscious awareness of viewpoint-dependent properties might arise. Response times were too slow to discriminate mere-sensory, perceptual, and post-perceptual processing stages. Finally, just as retinally projected shape does all the work that perspectival shape might be thought to do in explaining percept-formation (see *Scientific Reasoning*), retinally projected shape would appear to do all the work that perspectival shape might be thought to do in accounting for conscious appearance, conscious awareness, or conscious experience.

Much philosophical writing on consciousness relies heavily on introspection. This work, like work in philosophy of perception, is often scientifically uninformed. (For examples of philosophical writing on consciousness that makes intelligent use of science, see Block, 1995; 2011 and Phillips, 2021.) Within vision science and neuroscience, there are serious investigations of when consciousness arises during percept-formation, and of what brain areas support consciousness (Alais, Cass, O'Shea, & Blake, 2010; Brascamp, Blake, & Knapen, 2015; Dembski, Koch, & Pitts, 2021; Koch, Massimini, Boly, & Tononi, 2016; Lamme, 2010; Railo, Koivisto, & Revonsuo, 2011). Such work aligns with established science on percept-formation.

As noted, we are commonly aware of some elliptical shape corresponding to the projection cast by a rotated dinner plate. When people are asked to report retinally projected shape, they are often biased, reporting compromises between retinal and distal shape (see data in D. J. Cohen & Jones, 2008; Thouless, 1931; 1932). (Similar findings result when subjects are asked to report retinally projected size (Murray et al., 2006; Perdreau & Cavanagh, 2011; Shepard, 1990).) With training, and perhaps talent, such reports can become unbiased. In these circumstances, one is consciously aware of the actual projected shape. That awareness is painter's perspective (Alberti, 1435; D. J. Cohen & Jones, 2008; Perdreau & Cavanagh, 2011).

What is going on when painters are aware of, and accurately paint, the elliptical retinal projection while seeing rotated circular dinner plates? We know from science that painters do not *perceive* anything elliptical in such circumstances. The retinally projected shape is merely sensed, but not perceived. The elliptical-shape-awareness depends causally on this mere sensing. It is an open empirical question at what processing stage—mere sensory, perceptual, or post-perceptual—this conscious awareness emerges and what psychological capacities it requires. We think it likely that conscious mere sensing of retinally projected shape occurs at the perceptual stage of processing. We think it likely to be an aspect of *how* percepts represent distal shapes. We also think it likely that conscious awareness of it does not essentially require conceptual or other cognitive capacities. These conjectures await targeted investigation.

Conclusion

Let us summarize the view of perspective in shape-perception supported by modern vision science. From the retinal image, the visual system computes perceptual estimates of distal properties like distal shape and 3D orientation relative to the viewer. Retinally projected shape is jointly determined by distal shape and its 3D orientation. Whereas distal shape and its 3D orientation are perceptually represented, the shape of the retinal-image projection is merely sensed. Viewpoint-dependent properties of the retinal image not only contribute to computing perceptual estimates of distal shape and its 3D orientation. Such properties, including retinally projected shape and size, also contribute to *how*—in a specific viewpoint-dependent way—distal shape and its 3D orientation are represented. When a rotated distal circle is perceived, the associated elliptical shape-awareness is supported by how the rotated distal circle is perceived. Neither retinally projected shape, nor any perspectival shape distinct from it, is perceived.

Finally, we review our primary critical points. First, because of poor experimental controls (see Fig. 1AB), Morales et al. (2020) does not show that perception, sensing, or awareness of 'perspectival shape per se', retinal shape, appearance similarities, or 'representational similari[ties]' of any kind, ground explanation of its experimental results. Second, all of its results are consistent with standard findings and views in vision science, notwithstanding the authors' misrepresentation of those views. Third, vision science already invokes retinally projected shape and its encodings in well-supported explanations of percept formation. Since experiments cannot establish a new role for perspectival shape in human vision, such a role would add nothing to science and should be rejected. Fourth, sensing and perceiving, and different types of representation, should be better differentiated terminologically, to clarify important scientific distinctions.

For decades, science has advanced a well-supported account of the basics of shape-perception. There has never been a need to hold that perspectival shape, no matter how conceived, is perceived. More generally, there has never been a need to invoke perspectival shape—conceived as distinct from factors already cited by science—in accounts of conscious awareness, sensing, or perception. The article does not demonstrate a need to do so now.

We share the authors' view that better communication would benefit philosophy, psychology, and other sciences. We agree that philosophical accounts of perception should be more informed by science. We hope that our article contributes to these aims.

Methods

To compute the discriminability of the reaction-time data in the easy and hard conditions, we took, for each subject, the difference of the mean response times in each condition and divided by the average standard deviation (i.e. the square-root of the average response-time variances)

$$D = \frac{m_H - m_E}{\sqrt{\frac{1}{2}(v_E + v_H)}} \qquad\qquad (1)$$

where m_E and m_H are the mean response times, and v_E and v_H are the response-time variances, in the easy and hard conditions, respectively. Then, we averaged the discriminability across subjects to get the response-time discriminability for each experiment. Note that Equation 1 corresponds exactly to the expression for d-prime from signal detection theory (Green & Swets, 1966).

To compute differences in sensitivity (i.e. d-prime) between the easy and hard conditions from the accuracy data, we first converted proportion-correct to d-prime in each condition, using standard expressions from signal detection theory, and then took the difference

$$\Delta d' = d'_E - d'_H = 2[\Phi^{-1}(P_E) - \Phi^{-1}(P_H)] \qquad\qquad (2)$$

where $\Phi^{-1}(\cdot)$ is the inverse cumulative normal function, and P_E and P_H are proportion-correct in the easy and hard conditions, respectively.

For each experiment, we plotted the discriminability of the reaction-time data against the differences in sensitivity computed from the accuracy data. At the observed performance levels, an approximately linear relationship between these two quantities (see Fig. 2C) is expected under the drift-diffusion model, a commonly used tool for characterizing how speed and accuracy are related (Palmer et al., 2005). The correlation between the quantities is strong and significant (r=0.91, p=3.0x10^{-4}).

Note that some individual subjects responded correctly to all trials in the easy—and less commonly, in the hard—condition. When performance levels are at or near ceiling levels, confident estimates of sensitivity are notoriously difficult to obtain from a small number of trials. For example, given perfect performance on an experimental sample of 50 trials (the approximate number of trials performed by each subject in each condition) the nominal sensitivity (i.e. d-prime value) is infinite, and the 95% confidence interval on percent correct ranges from 92.9% to 100%. Estimating the sensitivity difference for each individual subject and then averaging across subjects, while preferable when confident estimates of individual-subject sensitivity can be computed, is impractical and fraught with difficulties in these circumstances. Thus, we computed sensitivity differences from the aggregated proportion-correct across subjects in each condition of each experiment. Simulations show that, for binomially distributed response data, the utilized method for computing sensitivity is a reliable and accurate method for computing average sensitivity differences in the group data. To be certain that the reported relationship between the speed and accuracy results does not depend on this analytical choice, we also computed the discriminability of the reaction-time data from the aggregated response times across subjects in each condition of each experiment. The correlation between the discriminability

of the aggregated reaction-time data and the sensitivity differences of the aggregated accuracy data is also strong and significant (r=0.90, p=4.5x10^{-4}).

Note also that the average response times that we report are very slightly different from the average response times reported by Morales et al. (2020). In the hard condition of Exp. 1, for example, we calculate an average response time of 540.8ms, whereas Morales et al. calculate an average of 542ms. These small discrepancies are due to the fact that, in each condition, Morales et al. averaged the response times only across trials to which observers responded correctly. We instead averaged response times in each condition across all trials, both correct and incorrect. We did so because we ultimately compare discriminability of the response-time data in the easy and hard conditions (Eqn. 1) to sensitivity differences computed from proportion correct in the easy and hard conditions (Eqn. 2). One must analyze both correct and incorrect trials to compute proportion correct. For purposes of comparison, it is thus appropriate to analyze the response times of both correct and incorrect trials. Regardless, differences between the two approaches to computing average response times have little effect on our correlational analysis (see Fig. 2C).

Data from Morales et al. (2020) and the analyses of that data described in this manuscript are available at https://osf.io/2eqsu/. The analyses described in this manuscript were not pre-registered.

References

Adams, W. J., Elder, J. H., Graf, E. W., Leyland, J., Lugtigheid, A. J., & Muryy, A. (2016). The Southampton-York Natural Scenes (SYNS) dataset: Statistics of surface attitude. *Scientific Reports*, *6*, 35805. http://doi.org/10.1038/srep35805

Alais, D., Cass, J., O'Shea, R. P., & Blake, R. (2010). Visual sensitivity underlying changes in visual consciousness. *Current Biology*, *20*(15), 1362–1367. http://doi.org/10.1016/j.cub.2010.06.015

Alberti, L.B. (1435/1991). On painting. London: Penguin Books.

Albright, T. D., & Desimone, R. (1987). Local precision of visuotopic organization in the middle temporal area (MT) of the macaque. *Experimental Brain Research*, *65*(3), 582–592. http://doi.org/10.1007/BF00235981

Arcaro, M. J., McMains, S. A., Singer, B. D., & Kastner, S. (2009). Retinotopic organization of human ventral visual cortex. *Journal of Neuroscience*, *29*(34), 10638–10652. http://doi.org/10.1523/JNEUROSCI.2807-09.2009

Backus, B. T., Banks, M. S., van Ee, R., & Crowell, J. A. (1999). Horizontal and vertical disparity, eye position, and stereoscopic slant perception. *Vision Research*, *39*(6), 1143–1170.

Baker, B., Lansdell, B., & Kording, K. (2021). A philosophical understanding of representation for neuroscience. *arXiv, 2102.06592*.

Barack, D. L., & Gold, J. I. (2016). Temporal trade-offs in psychophysics. *Current Opinion in Neurobiology*, *37*, 121–125. http://doi.org/10.1016/j.conb.2016.01.015

Barack, D. L., & Krakauer, J. W. (2021). Two views on the cognitive brain. *Nature Reviews Neuroscience*, *22*(6), 359–371. http://doi.org/10.1038/s41583-021-00448-6

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*, 227–247.

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*(12), 567–575. http://doi.org/10.1016/j.tics.2011.11.001

Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America. a, Optics, Image Science, and Vision*, *14*(7), 1393–1411.

Brascamp, J., Blake, R., & Knapen, T. (2015). Negligible fronto-parietal BOLD activity accompanying unreportable switches in bistable perception. *Nature Neuroscience*, *18*(11), 1672–1678. http://doi.org/10.1038/nn.4130

Burge, J. (2020). Image-computable ideal observers for tasks with natural stimuli. *Annual Review of Vision Science*, *6*, 491–517. http://doi.org/10.1146/annurev-vision-030320-041134

Burge, J., & Geisler, W. S. (2011). Optimal defocus estimation in individual natural images. *Proceedings of the National Academy of Sciences*, *108*(40), 16849–16854. http://doi.org/10.1073/pnas.1108491108

Burge, J., & Geisler, W. S. (2014). Optimal disparity estimation in natural stereo images. *Journal of Vision*, *14*(2), 1–18. http://doi.org/10.1167/14.2.1

Burge, J., & Geisler, W. S. (2015). Optimal speed estimation in natural image movies predicts human performance. *Nature Communications*, *6*, 7900. http://doi.org/10.1038/ncomms8900

Burge, J., McCann, B. C., & Geisler, W. S. (2016). Estimating 3D tilt from local image cues in natural scenes. *Journal of Vision*, *16*(13), 2. http://doi.org/10.1167/16.13.2

Burge, T. (2010). Origins of objectivity. Oxford: Oxford University Press.

Burge, T. (2014a). Perception: Where mind begins. *Philosophy*, *89*, 385–403. http://doi.org/10.1017/S003181911400014X

Burge, T. (2014b). Reply to Block: Adaptation and the Upper Border of Perception. *Philosophy and Phenomenological Research*, *89*, 573–583. http://doi.org/10.1111/phpr.12136

Burge, T. (2014c). Reply to Rescorla and Peacocke: Perceptual content in light of perceptual constancies and biological constraints. *Philosophy and Phenomenological Research*, *88*(2), 485–501.

Burge, T. (2022). Perception: First form of mind. Oxford: Oxford University Press.

Chin, B. M., & Burge, J. (2020). Predicting the partition of behavioral variability in speed perception with naturalistic stimuli. *Journal of Neuroscience*, *40*(4), 864–879. http://doi.org/10.1523/JNEUROSCI.1904-19.2019

Cohen, D. J., & Jones, H. E. (2008). How shape constancy relates to drawing accuracy. *Psychology of Aesthetics, Creativity, and the Arts*, *2*, 8–19. http://doi.org/10.1037/1931-3896.2.1.8

Cohen, J. (2010). Perception and computation. *Philosophical Issues*, *20*, 96–124.

Cottaris, N. P., Wandell, B. A., Rieke, F., & Brainard, D. H. (2020). A computational observer model of spatial contrast sensitivity: Effects of photocurrent encoding, fixational eye movements, and inference engine. *Journal of Vision*, *20*(7), 17. http://doi.org/10.1167/jov.20.7.17

Dembski, C., Koch, C., & Pitts, M. (2021). Perceptual awareness negativity: A physiological correlate of sensory consciousness. *Trends in Cognitive Sciences*, *25*(8), 660–670. http://doi.org/10.1016/j.tics.2021.05.009

Denison, R. N., Block, N., & Samaha, J. (2021). What do models of visual perception tell us about visual phenomenology? In F. De Brigard & W. Sinnott-Armstrong (Eds.), *Neuroscience and Philosophy*. Cambridge, MA.

DiCarlo, J. J., & Maunsell, J. H. R. (2003). Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *Journal of Neurophysiology*, *89*(6), 3264–3278. http://doi.org/10.1152/jn.00358.2002

Erdogan, G., & Jacobs, R. A. (2017). Visual shape perception as Bayesian inference of 3D object-centered shape representations. *Psychological Review*, *124*(6), 740–761. http://doi.org/10.1037/rev0000086

Fleming, R. W., Holtmann-Rice, D., & Bulthoff, H. H. (2011). Estimation of 3D shape from image orientations. *Proceedings of the National Academy of Sciences*, *108*(51), 20438–20443. http://doi.org/10.1073/pnas.1114619109/-/DCSupplemental

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, *96*(2), 267–314.

Gillam, B. (1967). Changes in the direction of induced aniseikonic slant as a function of distance. *Vision Research*, *7*(9-10), 777–783.

Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics (Vol. 1). New York: Wiley & Sons.

Hill, C. S. (2016). Perceptual relativity. *Philosophical Topics*, *44*, 179–200. http://doi.org/10.2307/26529411

Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: optimal cue combination. *Journal of Vision*, *4*(12), 967–992. http://doi.org/10.1167/4.12.1

Hume, D (1748/2000). An enquiry concerning human understanding, Section 12. T. Beauchamp (Ed.), Oxford: Oxford University Press.

Jaini, P., & Burge, J. (2017). Linking normative models of natural tasks to descriptive models of neural response. *Journal of Vision, 17*(12), 1-26. http://doi.org/10.1167/17.12.16

Kay, K. N., Weiner, K. S., & Grill-Spector, K. (2015). Attention reduces spatial uncertainty in human ventral temporal cortex. *Current Biology*, *25*(5), 595–600. http://doi.org/10.1016/j.cub.2014.12.050

Kemp, M. (1990). The science of art: Optical themes in Western Art from Brunelleschi to Seurat. New Haven: Yale University Press.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, *55*, 271–304. http://doi.org/10.1146/annurev.psych.55.090902.142005

Kim, S., & Burge, J. (2018). The lawful imprecision of human surface tilt estimation in natural scenes. *eLife*, *7*. http://doi.org/10.7554/eLife.31448

Kim, S., & Burge, J. (2020). Natural scene statistics predict how humans pool information across space in surface tilt estimation. *PLoS Computational Biology*, *16*(6), e1007947–26. http://doi.org/10.1371/journal.pcbi.1007947

Knapen, T. (2021). Topographic connectivity reveals task-dependent retinotopic processing throughout the human brain. *Proceedings of the National Academy of Sciences*, *118*(2). http://doi.org/10.1073/pnas.2017032118

Knill, D. C. (1998). Ideal observer perturbation analysis reveals human strategies for inferring surface orientation from texture. *Vision Research*, *38*(17), 2635–2656.

Knill, D. C., & Kersten, D. (2004). Visuomotor sensitivity to visual information about surface orientation. *Journal of Neurophysiology*, *91*(3), 1350–1366. http://doi.org/10.1152/jn.00184.2003

Knill, D. C., & Richards, W. (Eds.). (1996). *Perception as Bayesian inference*. Cambridge University Press.

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, *17*(5), 307–321. http://doi.org/10.1038/nrn.2016.22

Kupers, E. R., Carrasco, M., & Winawer, J. (2019). Modeling visual performance differences "around" the visual field: A computational observer approach. *PLoS Computational Biology*, *15*(5), e1007063. http://doi.org/10.1371/journal.pcbi.1007063

Kuroda, T. (1971). Distance constancy: functional relationships between apparent distance and physical distance. *Psychologische Forschung*, *34*(3), 199–219. http://doi.org/10.1007/BF00424606

Kwon, O.-S., Tadin, D., & Knill, D. C. (2015). Unifying account of visual motion and position perception. *Proceedings of the National Academy of Sciences*, *112*(26), 8142–8147. http://doi.org/10.1073/pnas.1500361112

Lamme, V. A. F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, *1*(3), 204–220. http://doi.org/10.1080/17588921003731586

Land, M. F., & Nilsson, D. E. (2002). Animal eyes. Oxford: Oxford University Press.

Lande, K. J. (2018). The perspectival character of perception. *The Journal of Philosophy*, *115*(4), 187–214. http://doi.org/10.5840/jphil2018115413

Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *Journal of Neuroscience*, *26*(51), 13128–13142. http://doi.org/10.1523/JNEUROSCI.1657-06.2006

Le, R., Witthoft, N., Ben-Shachar, M., & Wandell, B. (2017). The field of view available to the ventral occipito-temporal reading circuitry. *Journal of Vision*, *17*(4), 6. http://doi.org/10.1167/17.4.6

Linton, P. (2021). Conflicting shape percepts explained by perception cognition distinction. *Proceedings of the National Academy of Sciences*, *118*(10). http://doi.org/10.1073/pnas.2024195118

Locke, J. (1690/1975). An essay concerning human understanding. Book 2. In P. H. Nidditch (Ed.), Oxford: Oxford University Press.

Mackey, W. E., Winawer, J., & Curtis, C. E. (2017). Visual field map clusters in human frontoparietal cortex. *eLife*, 1–23. http://doi.org/10.7554/eLife.22974.001

Malik, J., & Rosenholtz, R. (1997). Computing local surface orientation and shape from texture for curved surfaces. *International Journal of Computer Vision*.

Mamassian, P. (1997). Prehension of objects oriented in three-dimensional space. *Experimental Brain Research*, *114*(2), 235–245. http://doi.org/10.1007/pl00005632

Marr, D. (1982). Vision. New York: W H Freeman & Company.

McCann, B. C., Hayhoe, M. M., & Geisler, W. S. (2018). Contributions of monocular and binocular cues to distance discrimination in natural scenes. *Journal of Vision*, *18*(4), 1–15. http://doi.org/10.1167/18.4.12

Mooney, S. W. J., Marlow, P. J., & Anderson, B. L. (2019). The perception and misperception of optical defocus, shading, and shape. *eLife*, *8*, 1–23. http://doi.org/10.7554/eLife.48214.001

Morales, J., Bax, A., & Firestone, C. (2020). Sustained representation of perspectival shape. *Proceedings of the National Academy of Sciences*, *117*(26), 14873-14882. http://doi.org/10.1073/pnas.2000715117

Morales, J., Bax, A., & Firestone, C. (2021). Reply to Linton: Perspectival interference up close. *Proceedings of the National Academy of Sciences*, *118*(28). http://doi.org/10.1073/pnas.2025440118

Murray, S. O., Boyaci, H., & Kersten, D. (2006). The representation of perceived angular size in human primary visual cortex. *Nature Neuroscience*, *9*(3), 429–434. http://doi.org/10.1038/nn1641

Norman, D. A., & Wickelgren, W. A. (1969). Strength theory of decision rules and latency in retrieval from short-term memory. *Journal of Mathematical Psychology*, *6*, 192–208.

Ogle, K. N. (1950). Researches in binocular vision. Philadelphia, London: W B Saunders.

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 1–29. http://doi.org/10.1167/5.5.1

Perdreau, F., & Cavanagh, P. (2011). Do artists see their retinas? *Frontiers in Human Neuroscience*, *5*, 1–10. http://doi.org/10.3389/fnhum.2011.00171

Phillips, I. (2021). Blindsight is qualitatively degraded conscious vision. *Psychological Review*, *128*(3), 558–584. http://doi.org/10.1037/rev0000254

Putnam, H. (1962). The Analytic and the synthetic. In H. Feigl and G. Maxwell (Eds.), Minnesota Studies in the Philosophy of Science III (Minneapolis: University of Minnesota Press, 1962), reprinted in H. Putnam, *Mind, Language, and Reality: Philosophical Papers, Vol. 2* (Cambridge: Cambridge University Press, 1975).

Putnam, H. (1973). Explanation and reference. In G. Pearce and P. Maynard (Eds.) Conceptual Change (Dordrecht: D. Reide), reprinted in H. Putnam, *Mind, Language, and Reality: Philosophical Papers, Vol. 2* (Cambridge: Cambridge University Press, 1975).

Putnam, H. (1975). The refutation of conventionalism. In M. Munitz (Ed.) Semantics and Meaning (New York: New York University Press), reprinted in H. Putnam, *Mind, Language, and Reality: Philosophical Papers, Vol. 2* (Cambridge: Cambridge University Press, 1975).

Railo, H., Koivisto, M., & Revonsuo, A. (2011). Tracking the processes behind conscious perception: a review of event-related potential correlates of visual consciousness. *Consciousness and Cognition*, *20*(3), 972–983. http://doi.org/10.1016/j.concog.2011.03.019

Rescorla, M. (2014). Perceptual constancies and perceptual modes of presentation. *Philosophy and Phenomenological Research*, *88*, 468–476. http://doi.org/10.1111/phpr

Rosenberg, A., Cowan, N. J., & Angelaki, D. E. (2013). The visual representation of 3D object orientation in parietal cortex. *Journal of Neuroscience*, *33*(49), 19352–19361. http://doi.org/10.1523/JNEUROSCI.3174-13.2013

Schellenberg, S. (2008). The situation-dependency of perception. *The Journal of Philosophy*, *105*, 55–84.

Shepard, R. N. (1990). Mind sights. WH Freeman/Times Books/Henry Holt & Co.

Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences*, *13*(11), 488–495. http://doi.org/10.1016/j.tics.2009.08.005

Stevens, K. A. (1983). Slant-tilt: The visual encoding of surface orientation. *Biological Cybernetics*, *46*(3), 183–195. http://doi.org/10.1007/BF00336800

Storrs, K. R., Anderson, B. L., & Fleming, R. W. (2021). Unsupervised learning predicts human perception and misperception of gloss. *Nature Human Behaviour*. http://doi.org/10.1038/s41562-021-01097-6

Thouless, R. H. (1931). Phenomenal regression to the real object. I. *British Journal of Psychology*, *21*, 339–359.

Thouless, R. H. (1932). Individual differences in phenomenal regression. *British Journal of Psychology*, *22*, 216–241.

Todd, J. T. (2004). The visual perception of 3D shape. *Trends Cogn Sci*, *8*(3), 115–121. http://doi.org/10.1016/j.tics.2004.01.006

Uchida, N., Kepecs, A., & Mainen, Z. F. (2006). Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nature Reviews Neuroscience*, *7*(6), 485–491. http://doi.org/10.1038/nrn1933

Vishwanath, D., Girshick, A. R., & Banks, M. S. (2005). Why pictures look right when viewed from the wrong place. *Nature Neuroscience*, *8*(10), 1401–1410. http://doi.org/10.1038/nn1553

von Helmholtz, H. (1909/1925). Helmholtz's treatise on physiological optics. Vol. III. Shouthhall, J. P. C. (Ed.) New York: Optical Society of America.

Wandell, B. A., & Winawer, J. (2011). Imaging retinotopic maps in the human brain. *Vision Research*, *51*(7), 718–737. http://doi.org/10.1016/j.visres.2010.08.004

Watson, A. B. (1979). Probability summation over time. *Vision Research*, *19*(5), 515–522. http://doi.org/https://doi.org/10.1016/0042-6989(79)90136-6

Welchman, A. E., Deubelius, A., Conrad, V., Bülthoff, H. H., & Kourtzi, Z. (2005). 3D shape perception from combined depth cues in human visual cortex. *Nature Neuroscience*, *8*(6), 820–827. http://doi.org/10.1038/nn1461

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, *41*, 67–85.