

Newly Sighted Perceivers and the Relation Between Sight and Touch

Penultimate draft (forthcoming, *Philosophy and the Mind Sciences*)

Abstract: Molyneux’s question asks whether a person born blind who has learned to identify shapes tactually could, if suddenly granted sight, immediately identify shapes visually. This question has been used to structure discussions of whether there is a “rational connection” between sight and touch—whether one may rationally doubt whether the same spatial properties are both seen and felt. I distinguish two issues under this general heading. The first concerns whether the visual and haptic perception of shape is rationally connected in *normally sighted perceivers*. The second concerns whether the visual and haptic perception of shape is rationally connected in *all possible perceivers*. I argue that real-world implementations of Molyneux’s question are irrelevant to the first issue, but potentially relevant to the second. However, I also argue that the second issue does not have the philosophical significance it is sometimes claimed to have. In particular, it cannot be used to adjudicate debates about whether the phenomenal character of perceptual experience is determined “externally,” by the worldly properties we perceive, or “internally,” by physiological or functional properties of our brains.

1. Introduction

Human perceivers have a remarkable ability to reidentify shape properties across sensory modalities. Shapes that are first encountered through touch can later be recognized by sight, and vice versa. Studies indicate that such cross-modal recognition reaches impressive levels of accuracy, sometimes even rivaling recognition performance within a modality (Desmarais et al., 2017; Norman et al., 2004). But what is the nature of the perceptual states mediating this capacity for cross-modal recognition? Does transfer of shape recognition across modalities depend on some intrinsic similarity between the perceptual states through which we apprehend shape in sight and touch, or are these states wholly distinct intrinsically but linked by mere association, as Berkeley famously claimed?¹

Traditionally, philosophers have examined the relation between visual and haptic shape perception through the lens of Molyneux’s question. In a pair of letters to John Locke in 1688 and 1692, the Dublin politician William Molyneux asked whether a person born blind who had learned to

¹ See Berkeley (1709, sects. 127, 147). Berkeley claims that “the extension, figures and motions perceived by sight are specifically different from the Ideas of touch called by the same names, nor is there any such thing as one Idea or kind of Idea common to both senses” (Berkeley, 1709, sect. 127). This view has come to be known as the *Heterogeneity Thesis*, and the claim that Molyneux’s question should be answered negatively forms a key premise in Berkeley’s argument for this thesis (van Cleve, 2007, p. 256). See Copenhaver (2014) for a nuanced discussion of Berkeley’s views on these issues, and see Prinz (2002, ch. 5) for a contemporary descendant of Berkeley’s Heterogeneity Thesis.

identify shapes by touch could, if suddenly granted sight, identify shapes visually (Locke, 1694, bk. II, ch. IX, sect. 8). Molyneux's question was long treated primarily as a thought experiment (Degenaar, 1996, p. 25), with occasional exceptions (Cheselden, 1728; Gregory & Wallace, 1963; Sacks, 1995). However, recent years have witnessed a surge of research into real-world cases of sight restoration, including one widely publicized attempt to resolve Molyneux's question empirically (Held et al., 2011). Theorists have vigorously debated the proper interpretation of these empirical results, and how best to devise an experimental test of Molyneux's question (Cheng, 2015; Connolly, 2013; Green, 2022a; Levin, 2018; Schwenkler, 2012, 2013, 2019). The present paper considers what can be learned from empirical investigations of newly sighted perceivers, and in particular how real-world implementations of Molyneux's question might bear on the theoretical issues animating the question.

We should differentiate Molyneux's question itself, which poses a concrete empirical test of a newly sighted subject, from various theoretical issues to which the question's answer might be relevant (see also Glenney, 2013; Hopkins, 2005; Levin, 2018; Matthen & Cohen, 2020). It is important to cleanly distinguish the latter issues, since it could turn out that empirical implementations of Molyneux's question are relevant to some of them but not others. Moreover, getting clear on the relation between Molyneux's question and the issues of more fundamental significance allows us to gauge whether Molyneux's question even matters today (Hopkins, 2005, p. 441). If, for instance, the only issues to which the question is relevant are ones that have already been empirically resolved, or for which we now have better lines of experimental attack, then the question is of merely historical interest, and not useful for guiding contemporary investigations of cross-modal perception.

One topic concerns whether there is a "rational" connection between the visual and haptic states through which we apprehend shape—i.e., whether it would be transparent to any sufficiently reflective subject that these states present the same worldly properties. Suppose that Jane first sees a square surface and then touches it while blindfolded, and perceives the surface veridically in both

cases. If the visual and haptic states through which Jane perceives squareness are rationally connected, then (assuming adequate attention, time, and motivation) she could not coherently doubt that she is perceiving the same shape property through both vision and touch, just as she could not coherently doubt that two visible squares seen side by side under optimal conditions appear the same (or approximately the same) shape.² Accordingly, cross-modal identification of visually and haptically perceived shape properties would be *uninformative* to the sufficiently reflective subject.

I take this topic to be the central focus of Gareth Evans’s classic paper, “Molyneux’s Question” (Evans, 1985). Evans presents two characters, V and B, who disagree concerning whether vision and touch support distinct systems of shape concepts. V holds that any subject who learns to apply a concept of some shape (say, *square*) on the basis of touch would be immediately prepared, on rational grounds, to apply the same concept on the basis of sight, and vice versa. Thus, there is just one system of shape concepts deployed flexibly on the basis of either sight or touch. B disagrees, holding that there is nothing intrinsic to the visual or haptic representation of *square* that prevents a subject from forming distinct, rationally unrelated concepts of *square* via the two modalities. If V is correct (as Evans seems to have thought), then it is presumably because vision and touch represent *square* in ways that leave no room for rational doubt about whether they represent the same property. Accordingly, it is transparent to the subject that the same shape concept is appropriately applied based on either representation.³ Conversely, B denies that this is the case. For B, it is *informative* for a subject to learn that her visual and haptic representations of *square* represent the same property.

² Of course, Jane might coherently doubt whether the objects *really are* the same shape: She might contemplate the possibility that one is seen illusorily while the other is not. Still, however, she would have no basis for doubting that they *visually appear* the same shape.

³ Evans’s argument for this view relies on the premises that vision and touch both represent shape properties within an egocentric framework, and that this framework is shared across modalities thanks to their common role in guiding behavior. Accordingly, vision and touch represent shape properties in the same way, or “in the same vocabulary” (Evans, 1985, p. 340)—i.e., as laid out thus-and-so within egocentric space. Thus, both modalities provide the same basis for applying shape concepts, and such concepts should be applied flexibly on the basis of either sight or touch. For discussion and critique of this argument, see Campbell (2005), Hopkins (2005), and Levin (2008).

Evans rightly observes that success in Molyneux's task would be insufficient to resolve the dispute between V and B (Evans, 1985, p. 330). For even if the newly sighted subject succeeds in identifying spheres and cubes by sight, this does not mean that her visual and haptic shape representations are rationally connected. Instead, there may be brute, hardwired associations that lead visual representations of a given shape to activate distinct haptic representations of that shape, and vice versa. The representations might be associated in this way without being rationally connected. Suppose, for instance, that when the newly sighted subject first sees a square surface, haptic imagery of squares spontaneously pops into mind, enabling her to accurately guess the surface's shape. Despite this, her visual and haptic representations of squareness might have no rational connection, since it may be coherent for her to doubt, upon careful consideration, whether they really denote the same property.⁴ Nonetheless, Evans thought that *failure* in Molyneux's task would be more decisive. If a newly sighted person could see well enough to visually represent the shapes of things around them, but nonetheless failed to recognize these shapes by sight (despite adequate time for reflection), then according to Evans, this result would "undermine V's position" (Evans, 1985, p. 330).

Notice that V may be correct that visual and haptic states are rationally connected even if it is possible for a subject to harbor various *non-rational* doubts about whether the states present the same shape properties. For example, a subject might fail to adequately attend to those aspects of the contents of the two states on which the rational connection depends, instead being distracted by various irrelevant differences between them. As Thomas Reid pointed out, the fact that visual shape perception is virtually always associated with perception of color, while haptic shape perception is virtually always associated with perception of hardness, may distract us from core commonalities in

⁴ Alternatively, visual and haptic shape representations may bear some "structural correspondence" that enables the newly sighted subject to make a reasonable inference about which object is the sphere and which is the cube, as Leibniz conjectured (*New Essays*, Bk. II, Ch. IX). Such structural correspondence might obtain without a genuinely rational connection between the representations Green, 2022a, p. 696).

their geometrical contents (Reid, 1764, ch. 6, sect. 7). The fundamental question, as I'll understand it, is not whether a subject can *ever* doubt whether their visual and haptic states represent the same shape properties, but whether a subject could *continue* to doubt this despite (i) having at least human-like reasoning capacities,⁵ (ii) engaging in adequate thought and reflection, (iii) attending selectively to the geometrical contents of the two states, and (iv) without committing any errors in reasoning. I'll describe subjects meeting conditions (i)-(iv) as "fully reflective." Thus, rational connection is partly a matter of perception, and partly a matter of cognition. For a subject's visual and haptic shape representations to be rationally connected, the perceptual representations must have contents or formats (perhaps *inter alia*) that are sufficiently similar that any fully reflective subject would be in position to cognitively determine, beyond any rational doubt, that they represent the same properties.

This paper considers the bearing of Molyneux's question on the topic of whether visual and haptic shape representations are rationally connected. However, two distinct issues fall under this heading. The first concerns whether the visual and haptic states that *normally sighted perceivers* employ to reidentify shape across modalities are rationally connected—i.e., whether any fully reflective perceiver with access to those very representations would be in position to determine beyond any rational doubt that they represent the same properties. Call this *Issue 1*:

Issue 1: Is there a rational connection between the particular visual and haptic representations of shape that normally sighted perceivers employ to reidentify shape properties across modalities?

However, even if visual and haptic representations of shape turn out to be rationally connected in normally sighted human perceivers, one might ask whether this holds for *all possible* perceivers who

⁵ This restriction is needed because there may be primitive creatures who perceive shape both visually and haptically in much the way we do, but have *no* reasoning capacities at all, and thus are unable to appreciate that their visual and haptic states represent the same properties. The mere possibility of such creatures should not suffice to show that our visual and haptic representations of shape are not rationally connected.

apprehend shape both visually and haptically.⁶ Thus, the second issue concerns whether visual and haptic representations of shape are *necessarily* rationally connected, or whether it is possible for a fully reflective subject to perceive the same shape through sight and touch while coherently doubting that this is the case. Call this *Issue 2*:

Issue 2: Are visual and haptic presentations of a given shape property *necessarily* rationally connected? That is, is it *impossible* for a fully reflective subject to enjoy visual and haptic states that present the same shape property while coherently doubting that this is so?⁷

Obviously, Issue 2 might be resolved negatively even if Issue 1 is resolved affirmatively.

I will not attempt to decisively settle Issues 1 or 2 here. Rather, my primary interests concern how studies of newly sighted perceivers may bear on each issue, and the issues' broader theoretical significance. Specifically, I will argue that investigations of newly sighted subjects are not helpful in settling Issue 1, but they may be helpful in settling Issue 2. Nonetheless, I'll also argue that Issue 2 does not have the philosophical significance that it is sometimes claimed to have.

The plan is as follows. Section 2 argues that empirical implementations of Molyneux's question are largely irrelevant to Issue 1. This is because any inference from newly sighted subjects' recognition performance to the relationship between visual and haptic shape representations in normally sighted subjects faces three serious pitfalls. Section 3 argues that data from newly sighted subjects may, however, be helpful in settling Issue 2—whether it is *possible* for a fully reflective subject to see and touch the same shape while coherently doubting whether this is so. Section 4 argues that the viability of “radically externalist” theories of spatial experience, namely those on which the phenomenal

⁶ Of course, one might explore intermediate questions as well: for instance, whether visual and haptic shape representations are rationally connected in all *human* perceivers (even if not in all possible perceivers), or whether they are rationally connected in all *conscious* perceivers. Thanks to [redacted] for this point.

⁷ I use the term “presentation” in posing this question to remain neutral on the precise relation between perceptual states and the worldly properties to which they relate the subject—e.g., whether this relation is best understood as a species of representation or instead as some non-representational relation like “acquaintance”—see section 3.2. Thus, theorists skeptical of the notion of perceptual representation may contemplate a version of Issue 2. Since I do not harbor such skepticism, I will set aside this complication until section 3, when the nature of perceptual shape experience takes center stage.

character of spatial experience is fully determined by the mind-independent spatial properties that we perceive, does not hinge on the resolution to Issue 2. While some take the mere possibility of rationally doubting whether vision and touch present the same shape properties to threaten externalism, the apparent threat rests, I contend, on a failure to appreciate the compositional character of shape perception. Section 5 concludes.

2. Issue 1: Pitfalls in extrapolating from the newly sighted

2.1. The Match Principle

This section considers what newly sighted perceivers can reveal about the perceptual representations mediating abilities for cross-modal recognition in the normally sighted. Would failure in Molyneux's test show that the shape representations responsible for cross-modal recognition in normally sighted perceivers are not rationally connected—i.e., that it would be possible for a fully reflective subject with access to both representations to coherently doubt that they represent the same property?

It is unclear whether Molyneux himself intended to raise this issue. As Degenaar (1996) observes, Molyneux may have been primarily concerned with whether three-dimensionality can be perceived by sight without prior association with touch.⁸ Nevertheless, regardless of the issue Molyneux *intended* to raise, many philosophers have taken his question to bear on whether there is a rational connection between the visual and haptic representations of shape that enable normally sighted perceivers to reidentify shapes across modalities. For instance, Schwenkler (2012) writes:

[T]hink of the way that shape and other spatial properties are perceived in sight and touch, of how we can tell right away whether a seen shape is the same as some felt one. Molyneux's question asks: can we do this only because of learned associations built up in the course of past experience, or are the representations of these properties related somehow intrinsically?
(186)

⁸ Likewise, Locke may have answered “no” to Molyneux's question simply because he thought three-dimensional spatial properties were imperceptible by sight without prior association with touch, or acts of judgment (Bruno & Mandelbaum, 2010).

This section argues that newly sighted subjects' perceptual capacities are largely irrelevant to the relationship between visual and haptic shape representations in the normally sighted, so Molyneux's test cannot be used for the purpose Schwenkler recommends. That is, the test cannot be used to shed light on how most of us (i.e., mature, normally sighted perceivers) can "tell right away whether a seen shape is the same as some felt one." My discussion emphasizes empirically documented differences in spatial perception between normally sighted and newly sighted (and congenitally blind) perceivers. It extends prior observations by Green (2022a) to highlight two further pitfalls in generalizing from newly sighted subjects' recognitional capacities to the relation between visual and haptic representations in the normally sighted.

First, however, we should understand *why* differences in either visual or haptic shape representation between normally and newly sighted subjects threatens the evidential value of Molyneux's test vis-à-vis Issue 1. Figure 1 illustrates the problem. Molyneux's question proposes a test of newly sighted subjects' recognitional capacities: Are their visual and haptic representations of shape related in such a way that shapes previously recognizable through touch can be immediately reidentified by sight? Suppose that we run the experiment and the results clearly support either a positive or a negative answer. For this outcome to bear evidentially on the relation between visual and haptic shape representations in *normally* sighted subjects, it must be that the relation between newly sighted subjects' visual and haptic shape representations is approximately preserved, or *matched*, in normally sighted subjects. And to be confident that *this* is true, we must be confident that visual and haptic shape representations are relevantly alike in the two groups of subjects. Green (2022a) calls this the *Match Principle*:

Match Principle: Newly sighted perceivers form visual and [haptic] representations of shape that are intrinsically similar to the visual and [haptic] representations of shape directly responsible for cross-modal shape recognition in normally sighted perceivers. (Green, 2022a, p. 698)

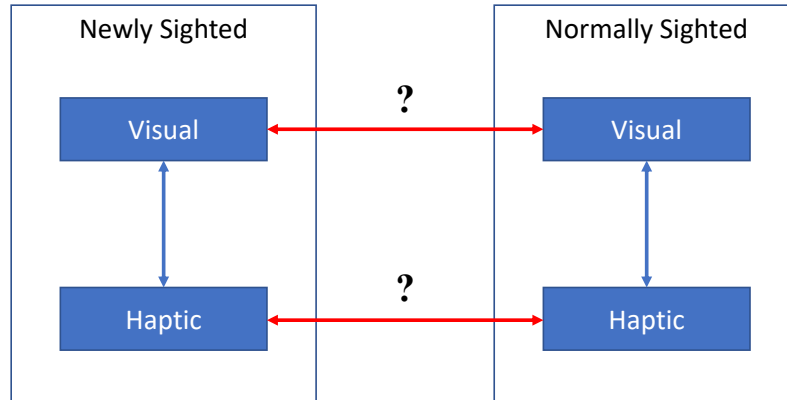


Figure 1. Illustrating the Match Principle. Molyneux’s test plausibly reveals something about the relationship between visual and haptic shape representations in *newly* sighted subjects. But for the test to bear evidentially on the relationship between visual and haptic shape representations in *normally* sighted subjects, we must be confident that visual and haptic shape representations are relevantly alike in the two cases. Otherwise, we have no reason to believe that the relation between newly sighted subjects’ visual and haptic shape representations is preserved in normally sighted subjects.

The notion of “intrinsic similarity” is admittedly vague. However, it is meant to call attention to similarities in features that the representations possess aside from any history of cross-modal association. For example, newly sighted and normally sighted subjects might form visual and haptic representations that share a common format, content, or neural substrate, but differ only in whether associations between those representations have been established.⁹ If so, then the Match Principle would hold. Conversely, if newly sighted and normally sighted subjects form visual or haptic shape representations which differ markedly in format, content, or neural substrate, then the Match Principle may fail to hold.

To see why the Match Principle is important, suppose that newly sighted subjects fail Molyneux’s test. As noted earlier, Evans takes this result to undermine the view that our visual and haptic shape representations are rationally connected. However, if the Match Principle is false, then

⁹ See Di Stefano and Spence (2023) for helpful discussion of various potential forms of intrinsic similarity between representations from different sense modalities. Note that representations from vision and touch might be similar without sharing precisely the same content or format. For instance, they might exhibit some sort of structural isomorphism (Di Stefano & Spence, 2023, pp. 10-12).

the following possibility arises: Visual and haptic shape representations *are* rationally connected in normally sighted perceivers (i.e., any fully reflective subject who had access to *those very representations* would be able to determine that they represent the same shape properties), but they are not rationally connected in newly sighted perceivers, due to differences in content or format between the two. In other words, the intrinsic similarities (e.g., similarities in content or format) between vision and touch that underlie a rational connection within normally sighted perceivers might be missing in newly sighted perceivers.

This section argues that the Match Principle confronts three serious difficulties that undercut the evidential relevance of Molyneux's test to Issue 1. Newly sighted subjects simply can't reveal much about the relation between visual and haptic shape representation in normally sighted subjects, so they should not be relied on for this purpose.

2.2. *The convergent processing problem*

The first difficulty for the Match Principle is what I'll call the *convergent processing problem*. The problem can be illustrated as follows. Suppose that in normally sighted perceivers, shape is represented differently at different levels of processing. For simplicity, assume that there are two visual shape representations, V1 and V2, and that the construction of V1 precedes and causes the construction of V2. Assume that a similar situation obtains for touch: construction of representation T1 precedes and causes the construction of T2. Finally, suppose that there is a rational connection between V2 and T2, but *not* between either V1 and T1, V1 and T2, or T1 and V2. If so, then the visual and haptic processing of shape *converges* to a rational connection. Visual and tactual shape representations are rationally *disconnected* at earlier processing stages, but rationally connected at later stages.

Crucially, if cross-modal recognition is mediated by V2 and T2, then Issue 1 is settled affirmatively: The visual and haptic shape representations that mediate cross-modal recognition in

normally sighted subjects *are* rationally connected. However, suppose that the higher-level visual processes responsible for forming V2 are impaired in newly sighted subjects. Then it is possible that such perceivers are capable of forming V1, but not V2. If so, then even though there *is* a rational connection between visual and haptic shape representation in normally sighted perceivers, this connection is absent in newly sighted perceivers. Accordingly, failures of cross-modal recognition in newly sighted subjects would not impugn a positive resolution to Issue 1.

So far, the convergent processing problem is merely hypothetical. Should we be concerned that it is actual? Yes. The evidence indicates that cross-modal shape recognition in normally sighted perceivers relies on part-based, view-invariant shape representations formed at later stages of sensory processing, both in vision and in touch. However, there is also evidence that mid- and high-level visual processing is significantly disrupted in newly sighted perceivers. Thus, there is a realistic possibility that the visual and haptic representations that mediate cross-modal recognition in normally sighted perceivers are rationally connected, but this connection is missing in newly sighted perceivers because they fail to form the relevant visual representations.

Various studies have examined our ability to recognize previously touched shapes on the basis of sight, or vice versa (Desmarais et al., 2017; Lawson, 2009; Norman et al., 2004; Yildirim & Jacobs, 2013; see Lacey & Sathian, 2014 for a review). But what kinds of representations explain this ability? Note that there are myriad schemes that a perceptual system might use to represent shape (Elder, 2018; Green, 2023; Hummel, 2013; Lande, 2024). However, one significant distinction is between *view-invariant* schemes that encode shape in roughly the same way across certain changes in orientation, and *view-dependent* schemes that encode shape differently across most changes in orientation.

No contemporary scheme of perceptual shape representation is *fully* view-invariant (Biederman & Gerhardstein, 1993; Hummel, 2013). If an object rotates to bring previously unseen (or unfelt) parts into view, then the perceptual representation of its shape will often change. Furthermore,

even smaller changes in orientation should be expected to alter the *precision* with which various parts of a shape are represented (Erdogan & Jacobs, 2017). View-invariance versus view-dependence is really a matter of degree. However, where a representation scheme falls on the spectrum between view-dependent and view-invariant is determined largely by the reference frame it employs. *Object-centered* schemes encode shape relative to an object's intrinsic axes, such as medial axes, axes of symmetry, or axes of elongation (Chaisilprungraung et al., 2019; Feldman & Singh, 2006; Green, 2023; Quinlan & Humphreys, 1993; Spelke, 2022). *Viewer-centered* schemes encode shape relative to axes that are privileged for the viewer. For instance, shape might be coded as a vector of critical features within a reference frame centered on the hands or the eye, with axes fixed by the structure of the hand (touch), or the structure of the retina (vision), or perhaps the direction of gravity (Poggio & Edelman, 1990; Ullman & Basri, 1991). There are also mixed schemes, in which certain aspects of shape are encoded in a viewer-centered manner and other aspects in an object-centered manner. For example, Biederman's (1987) geon theory encodes the shapes of parts of an object in a fully object-centered manner, but characterizes relations between parts (e.g., *above*) relative to the viewer or direction of gravity.

Importantly, when an object rotates with respect to the viewer, this alters its spatial layout within a viewer-centered reference frame. Thus, viewer-centered representations encode the same shape differently (e.g., via a different series of coordinates) at different orientations. Conversely, object-centered representations can remain approximately invariant across certain changes in orientation (modulo the qualifications above) because the layout of an object's boundaries relative to its intrinsic axes remains stable across such changes, assuming the object moves rigidly.

In a key study, Lacey et al. (2007) examined patterns of view-dependence for both within-modal and cross-modal shape recognition. Shape recognition is view-dependent when it is superior (as regards accuracy or reaction time) when a shape is presented at the same orientation ("view") as it

possessed during initial familiarization with the shape. Lacey et al. had participants first study a set of four objects made up of rectangular blocks in varying spatial arrangements. The study phase was either purely visual or purely haptic. Next, the participants were tested on a single object, where the task was to report which of the initial four objects it matched. The test phase could be either visual or haptic, and the object was either presented at the same orientation as the study phase or was rotated 180°. This allowed the authors to test for view-dependence both within and across modalities. Consistent with typical patterns (Newell et al., 2001; Peissig & Tarr, 2007), Lacey et al. found that within-modal shape recognition was view-dependent. In both vision and touch, participants were more accurate when shapes were tested at the same orientation as they had been presented during the study phase. Cross-modal recognition, however, was *view-invariant*. When shapes were tested in a different modality than they were studied, recognition was equally accurate regardless of orientation.¹⁰

Lacey and colleagues take these findings to suggest that vision and touch initially produce modality-specific representations that encode an object's shape in a viewer-centered manner. Because processes of within-modal shape recognition depend at least partially on these earlier modality-specific representations, within-modal recognition is superior at familiar viewpoints.

However, the fact that cross-modal recognition was view-invariant suggests that at later stages of processing, both vision and touch may generate object-centered representations of shape that remain stable across certain changes in viewpoint (Erdogan et al., 2015; Green, 2022a, 2022b). If cross-modal recognition is primarily mediated by these object-centered representations, then we have a natural explanation of why cross-modal recognition is equally accurate at familiar and unfamiliar orientations. Physiological support for this model comes from studies showing that both visual and haptic shape perception activate the lateral occipital complex (LOC), a high-level sensory area selectively responsive to global shape (Amedi et al., 2001, 2010; Erdogan et al., 2016; Masson et al.,

¹⁰ These results have since been replicated (Lacey et al., 2010; Ueda & Saiki, 2012).

2016), which has also been found to display some degree of view-invariance in shape representation (James et al., 2002; Lescroart & Biederman, 2013).

Thus, the evidence suggests that the perceptual representations mediating cross-modal recognition in normally sighted perceivers have a strong degree of view-invariance, and are formed in later stages of perceptual processing (potentially LOC). Cross-modal recognition does not seem to rely on the earlier view-dependent representations usable for within-modal recognition. Crucially, however, research on *newly* sighted perceivers suggests that their mid- and high-level visual processes are significantly impaired relative to normal perceivers, casting doubt on their ability to form view-invariant shape representations. I will only discuss a small selection of findings here; see May et al. (2022) for a broader review.

In a study of subject M.M., who lost vision at age three and had sight restored forty years later, Fine et al. (2003) documented a number of striking deficits. M.M. was able to name simple 2D shapes, but was severely impaired in perceiving 3D shapes in perspective line drawings, shapes defined by subjective contours, and overlapping transparent surfaces.¹¹ He was able to recognize only 25% of common objects, and although he could discriminate faces from non-faces, he could not reliably classify them by gender or facial expression. In a follow-up study over a decade later, Huber et al. (2015) found no significant improvements in M.M.'s visual capacities. They also tested other abilities, including the matching of 3D shapes at different orientations in depth. M.M.'s performance in this task was statistically indistinguishable from chance. Thus, higher-level visual functions—including, critically, the ability to visually represent shape in a manner that supports reidentification across changes in viewpoint—appear to be significantly impaired in M.M.

¹¹ See also Putzar et al. (2007) for evidence that subjects who experience visual deprivation during the first several months of life struggle to perceive illusory contours in Kanizsa figures. And see Ostrovsky et al. (2009) for evidence that newly sighted subjects have difficulty parsing scenes into objects on the basis of static Gestalt cues.

McKyton et al. (2015) confirmed this general pattern of findings in a group of 11 sight-recovery subjects. The subjects were shown displays of 4-12 elements and asked to detect the “oddball” item that differed from the rest. While the newly sighted subjects could reliably identify targets that differed from distractors in simple features like color, size, or 2D shape, they were at chance in identifying targets that differed in 3D shape or 2D shape defined by modal or amodal completion. Other evidence suggests that newly sighted subjects are also impaired in temporally integrating information about local parts of an object to form a global shape representation (Orlov et al., 2021).

These behavioral results suggest that newly sighted perceivers struggle to visually represent the sorts of properties typically recovered during mid- and high-level visual processing. Physiological results bolster this conclusion. Through EEG recording, Sourav et al. (2018) monitored the event-related potentials evoked by visible line gratings in subjects who had congenital cataracts removed later in life. Visual stimuli commonly elicit electrophysiological responses with two components: the C1 wave, which begins about 50 ms after stimulus onset and is thought to reflect processing in primary visual cortex (V1), and the P1 wave, which follows the C1 wave and is believed to reflect activity in downstream, extrastriate areas (Woodman, 2010). Importantly, Sourev et al. found that subjects with congenital cataracts showed roughly typical C1 waves. However, the amplitude of the P1 component was substantially lower in the late-sighted subjects relative to controls, suggesting impaired downstream visual processing. Sourav et al. conclude: “[B]asic features of retinotopic processing in the early visual cortex [are] functional with a typical time-course after a period of bilateral congenital blindness, whereas extrastriate processing does not seem to recover to the same extent” (Sourav et al., 2018, p. 14). Likewise, in their review of newly sighted subjects’ visual capacities, May et al. (2022) emphasize the “persistent deficits in some extra-striate visual abilities after extended visual deprivation” (p. 12).

Such results comport with abundant evidence that the visual cortex develops atypically if visual input is absent for a prolonged period early in life (Noppeney et al., 2005; Wiesel & Hubel, 1963). Due to neuronal deterioration in the absence of visual input during early critical periods, it also cannot be assumed that newly sighted subjects are an adequate substitute for normally sighted newborns (see also Gallagher, 2005, pp. 165-167). In fact, there is evidence that cross-modal transfer of shape recognition is present to some degree within the first few days of life (Sann & Streri, 2007; Streri & Gentaz, 2003). However, the representations mediating such cross-modal recognition in newborns may well be missing in newly sighted subjects.

To sum up: The evidence suggests that cross-modal shape recognition in normal perceivers depends on highly view-invariant shape representations formed in later stages of perceptual processing (perhaps within LOC). However, mid- and high-level visual processing in newly sighted perceivers is seriously impaired. This raises the realistic possibility that even if the representations responsible for cross-modal recognition in normal perceivers *are* rationally connected, newly sighted perceivers may lack access to this connection because they are unable to form the representations needed to forge it.¹² So, newly sighted subjects should not be relied on to resolve Issue 1.

2.3. *The coordinated processing problem*

A second problem threatening the Match Principle arises from the fact that our sensory systems are malleable over time. In particular, a history of cross-modal interaction with vision may affect the way shape is *haptically* represented in mature, normally sighted perceivers. This possibility further

¹² Green (2022a, 2022b) argues further that the visual and haptic representations directly responsible for cross-modal shape recognition in normally sighted subjects are plausibly rationally connected—indeed, that the two are *type-identical*. I do not seek to establish this further claim here. My present point is only that this possibility is clearly left open even if newly sighted subjects fail Molyneux's test.

problematizes the extrapolation from newly sighted subjects' perceptual shape representations to those of normal perceivers.

O'Callaghan (2019) observes that auditory experience is liable to differ between a hypothetical subject who has only ever undergone auditory experience, and typical subjects, for whom auditory experience has regularly occurred alongside experiences in other modalities: "If you could only ever hear, but not see, touch, taste, or smell, then your auditory experience could differ from how it actually now is" (p. 128). Here I suggest that something analogous occurs between touch and vision. Haptic spatial representation differs between congenitally blind subjects and normally sighted subjects. The evidence suggests that touch learns from vision and vice versa during development, refining the systems of spatial representation employed in both modalities.¹³

Patterns of interaction across modalities are known to change during development. While adults integrate visual and haptic estimates of size and slant in a statistically optimal manner (Ernst & Banks, 2002), multiple studies suggest that children below 8 years of age often fail to do this, instead permitting one modality to fully dominate the other (Gori, 2015; Gori et al., 2008; Nardini et al., 2010). Might this development of multisensory interaction also influence the nature of the representations *within* modalities? If so, then haptic spatial representation may differ between congenitally blind and sighted subjects. Consistent with this, Gori et al. (2010) found that congenitally blind participants performed worse than sighted counterparts in haptic discrimination of orientation (see also Postma et al., 2008). Gori et al. suggest that haptic processing of orientation is refined and calibrated by vision,

¹³ More generally, haptic representation is plausibly sensitive to an individual's learning history. While the distinction between blind and sighted subjects is a particularly extreme example of this phenomenon, subtle differences also arise within the sighted population—e.g., between experts and novices in certain haptic tasks. For example, expert table-tennis players exhibit superior haptic discrimination of racket vibration and angle relative to novices (Park & Kim, 2014).

enabling the haptic system to correct for systematic biases, since vision possesses more reliable cues to this property (see also Gori, 2015).¹⁴

Other studies have investigated differences in haptic shape representation within congenitally blind individuals. One such difference concerns susceptibility to *prototype effects*. Normally sighted subjects are better at visually recognizing shapes when they are presented in canonical or prototypical orientations. For example, rectangles are better recognized when they are oriented horizontally or vertically (rather than diagonally), and triangles are better recognized when their bases are oriented horizontally (Kalenine et al., 2011). It is natural to wonder whether such prototype effects are mirrored in haptic shape recognition. Interestingly, Theurel et al. (2012) found that the answer differs for sighted and congenitally blind perceivers. Subjects were asked to identify squares, rectangles, or triangles by touch, discriminating them from distorted distractors. For blindfolded normally sighted subjects, haptic recognition was faster at canonical orientations, just as in vision. However, congenitally blind subjects were equally fast regardless of orientation. This suggests that the visual system represents shape in a manner that prioritizes canonical orientations (Humphreys & Quinlan, 1988; Spelke, 2022b, ch. 6; Tarr & Pinker, 1990), and that the haptic system learns to adopt this method as well, but only after coordination with vision.

Summing up, the Match Principle requires newly and normally sighted subjects to resemble not only in the *visual* representation of shape, but *haptic* representation of shape as well. However, there is evidence that haptic spatial representation differs in key respects between blind and sighted

¹⁴ Conversely, Gori et al. (2010) found that haptic size discrimination was superior in the blind participants, suggesting that, unlike orientation, touch does not need to be calibrated by vision for size perception (see also Gori et al., 2012). Likewise, haptic discrimination of 2D angles is superior in blind perceivers (Alary et al., 2008). Thus, while haptic representation of spatial features differs between blind and sighted populations, whether such features are represented with greater or lesser fidelity in blind perceivers plausibly varies depending on whether vision or touch has better cues to the relevant feature (Alary et al., 2009; Gori, 2015, p. 84).

adults. These differences raise the possibility that the Match Principle fails not only on the visual side, but the haptic side too.

2.4. *The contour classification problem*

So far, I have examined problems associated with extrapolating from newly sighted perceivers' recognition performance to conclusions about the relation between visual and haptic shape representations in normally sighted perceivers. Challenges to this inference arise because newly sighted perceivers are atypical in both their visual and haptic processing of shape. Thus, we should have serious doubts about the Match Principle. However, one might think that we can evade these concerns by testing recognition for simpler features like edges or corners.¹⁵ After all, the foregoing exceptions to the Match Principle arise most markedly in mid- to high-level vision, where the visual impairments of newly sighted subjects are most pronounced. Perhaps by focusing on properties thought to be recovered in "early" vision, we identify a domain where these problems for the Match Principle can be avoided. If newly sighted subjects' visual and haptic representations of edges and corners are sufficiently like ours, then we might take their abilities to visually identify such local features to indicate whether visual and haptic representations of these features are rationally connected in normally sighted perceivers.

While this is a tempting thought, there is reason to doubt that variants of the Match Principle obtain even for simple, one-dimensional features like edges. For although local luminance, texture, or color discontinuities are *detected* in the earliest stages of vision, edges or corners *as we perceive them* go well beyond the deliverances of such primitive feature detection. Mature perception of edges is shaped, *inter alia*, by higher-level processes of contour *classification*, which determine whether a given

¹⁵ Levin (2008, pp. 8–9) proposes testing newly sighted subjects on edges or corners, but for reasons distinct from the problems discussed in sections 2.2-2.3. Levin suggests that if the newly sighted perceiver can visually identify lines or edges, we can be more confident that she is not relying on guesswork or analogical inference to do so.

discontinuity in the retinal image corresponds to a material edge, cast shadow, change in illumination, or surface scratch (Kellman et al., 2001; Kellman & Fuchser, 2023). Such classification is based on various sources of information. For instance, material edges between surfaces are likely to be signaled by multiple cues together, such as discontinuities in motion, luminance, color, and binocular disparity, while other sorts of contours (e.g., shadow boundaries) are typically associated with only a narrow subset of these.

Unfortunately, processes of contour classification are highly impaired in newly sighted subjects. Ostrovsky et al. (2009) showed newly sighted subjects images of objects displaying variations in shading, texture, and partial occlusion. The participants were asked to enumerate the objects present in the image. Substantial errors occurred in this task. The participants treated all regions of uniform hue or luminance as separate objects. For example, the shaded region of a ball was treated as a separate object from the more brightly illuminated region (see figure 2). It seems, then, that newly sighted perceivers are largely insensitive to whether a given discontinuity in the retinal image corresponds to a material edge, cast shadow, or texture difference. It is natural to speculate that proper functioning of higher visual areas (and perhaps feedback to V1/V2) is critical to contour classification (e.g., Drewes et al., 2016; Mathes & Fahle, 2007), and this explains newly sighted subjects' deficits. But irrespective of the underlying mechanism, the important point is that newly sighted subjects likely do not *see* edges and corners the way normally sighted perceivers do.




Figure 2. The top row shows the boundaries drawn by a newly sighted patient (S.K.) when asked to divide the images into distinct objects. The bottom row shows the outputs of a simple algorithm that separated the images into regions based on similarities in hue and luminance, without sensitivity to contour classification—e.g., whether a given edge corresponds to a genuine material edge or only a change in surface color or illumination. Notice the similarities between S.K.’s strategy and the simple algorithm. *Source:* Ostrovsky et al. (2009). Reprinted with permission of Sage Publications.

Such deficits raise grounds for skepticism about the Match Principle even for perceptual representations of edges and corners. For, in normally sighted subjects, the matching of edges and corners across modalities plausibly relies on representations that successfully distinguish material edges from other sorts of discontinuities in the retinal image. After all, while material edges are perceptible through both sight and touch, discontinuities in color or illumination are typically perceptible only by sight. Contour classification determines which contours in the retinal image are *viable candidates* for cross-modal matching. Accordingly, even if newly sighted subjects failed to identify edges or corners by sight, we could not determine whether this failure should be attributed to a universal lack of rational connection between visual and haptic representations of these features, or instead to specific impairments in the visual representation of these features brought on by visual deprivation. Of course, newly sighted subjects might be able to *guess* that a given visible luminance discontinuity ought to be matched to a given tangible surface edge. But these representations may well

lack a rational connection in the absence of intact contour classification, since the visual representation remains non-committal about whether that discontinuity corresponds to a genuine material edge.

Thus, we should be skeptical of any extrapolation from newly sighted subjects' ability to visually identify edges or corners to the relation between visual and haptic representations of these features in normally sighted subjects. I conclude that the Match Principle is implausible not only in the case of global shape, but also in the case of simple contour features, since perception of these features is *also* atypical in newly sighted subjects.

3. Issue 2: Newly sighted subjects and the possibility of rational disconnect

So far, I've argued that newly sighted subjects are a poor source of evidence concerning Issue 1: whether the visual and haptic shape representations that normally sighted perceivers rely on for cross-modal recognition are rationally connected. They are a poor source of evidence because, due to empirically documented differences in spatial perception between newly and normally sighted perceivers, it is a realistic possibility that visual and haptic shape representations are rationally connected in normally sighted perceivers, but not in newly sighted perceivers.

I now turn to Issue 2: whether it is *necessary* that visual and haptic presentations of the same shape property are rationally connected. I argue that newly sighted subjects *may* bear on this issue, but whether they do depends on whether they genuinely represent *shape* in vision, or merely unorganized clusters of contour features.

To evaluate the relevance of newly sighted subjects to Issue 2, let's consider a recent, widely discussed attempt to settle Molyneux's Question empirically. Held et al. (2011) examined five newly sighted subjects, all within 48 hours after sight restoration surgery. The subjects were presented with shapes resembling Lego blocks, and their recognition was assessed under three conditions. First, in the *visual-visual* (V-V) condition, subjects saw a sample shape, then saw that shape again alongside a

distractor object, and had to select the shape they had seen earlier. In the *tactual-tactual* (T-T) condition, this process was repeated, except all the shapes were felt rather than seen. Finally, in the crucial *tactual-visual* condition (T-V), the shape was first felt, and then had to be reidentified visually alongside a distractor.

Remarkably, Held et al. found that performance was near ceiling in both the V-V and T-T conditions (92% and 98%, respectively), but near chance in the T-V condition (58%; see also Chen et al., 2016).¹⁶ This combination of results might be taken to indicate that: (i) Newly sighted subjects successfully represent shape in both sight and touch, accounting for strong performance in the V-V and T-T conditions, but (ii) their visual and haptic states represent shape in qualitatively different ways, permitting doubt about whether the same shape property is both seen and felt, and this accounts for poor performance in the T-V condition. Thus, it is possible for a fully reflective subject who perceives shape through both sight and touch to coherently doubt whether the same shape properties are presented through both modalities, and Issue 2 should be settled negatively.

Follow-up discussions of Held et al. (2011) have tended to focus on whether (i) is correct. In particular, Schwenkler (2012, 2013) hypothesizes that newly sighted subjects might have succeeded in the V-V condition not by matching representations of global shape, but by reidentifying local contour features (e.g., edges or corners) that differed between the target shape and the distractor.¹⁷ Crucially, if (i) is false, then the Held et al. results do not straightforwardly bear on Issue 2. Recall that Issue 2 asks whether it is possible for a perceiver to enjoy visual and haptic presentations of the same *shape property* that are rationally disconnected. If newly sighted subjects do not enjoy visual perception of shape, but only local contour features, then they are irrelevant to the issue.

¹⁶ Note, however, that performance in the T-V condition improved to over 80% only 5 days after surgery, suggesting that newly sighted perceivers may acquire basic cross-modal recognition abilities fairly rapidly.

¹⁷ In fact, Held et al. themselves propose that visual shape recognition in the newly sighted is best explained by “an account that relies on strategies using two-dimensional features, such as corners, edges and curved segments,” rather than by a “visual ability to create a three-dimensional shape representation” (Held et al., 2011, p. 552).

In light of Schwenkler’s concerns, theorists have made various suggestions about how to refine the stimuli or presentation conditions to make it easier for newly sighted subjects to form global shape representations (e.g., Cheng, 2015; Connolly, 2013; Levin, 2018; Schwenkler, 2012, 2013)—for instance, by using raised line-drawings rather than 3D shapes as stimuli, or by giving subjects the opportunity to explore the objects from multiple perspectives.

I don’t know whether any of these manipulations would generate improvements in the T-V condition. However, given that the primary source of uncertainty in interpreting newly sighted subjects’ performance in the Held et al. task concerns whether they are able to form visual representations of global shape, over and above local contour features, I suspect that stimuli specifically designed to investigate this question in other domains may be informative. In particular, parallels may be found in recent research on the visual capacities of deep neural networks.

To examine whether deep neural nets such as AlexNet classify images based on global shape or merely collections of local features, Baker et al. (2020) created stimuli that disentangled the two—for instance, a global square-like shape composed of local curved segments or a global circle-like shape composed of local straight segments (fig. 3). We immediately appreciate the similarity between such figures and “ideal” perfect squares and circles.¹⁸ Thus, visual shape representations are not like lists of contour features, since these lists would be very different for the stimuli in figure 3 and the ideal squares and circles that they resemble. Interestingly, however, networks that could reliably identify ordinary squares and circles were biased toward classifying the square-like images composed of curved segments as *circles*, suggesting marked impairment in the networks’ extraction of global shape information (see also Baker et al., 2018; Baker & Elder, 2022).

¹⁸ Recall Koffka’s famous remark that in form perception “the whole is something else than the sum of its parts” (Koffka, 1935., p. 176).

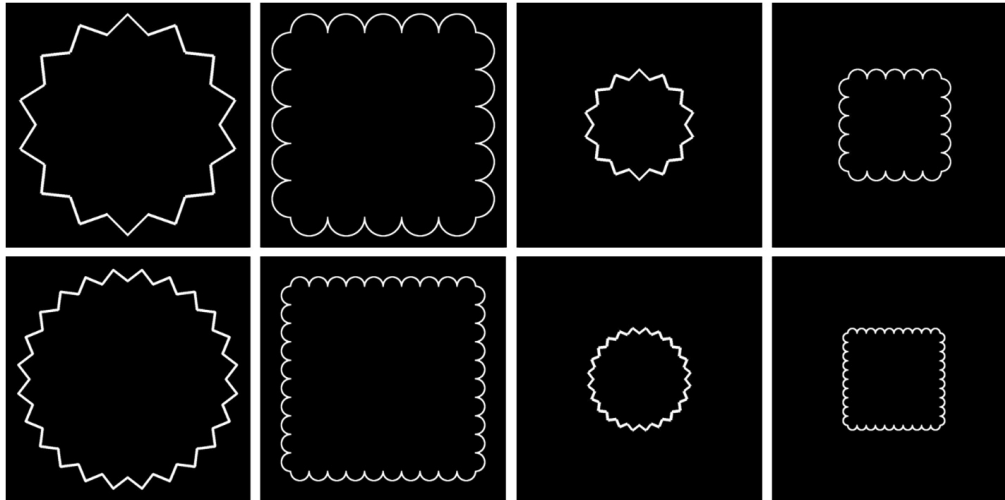


Figure 3. Stimuli used by Baker et al. (2020) to disentangle local contour features and global shape. *Source:* Baker et al. (2020). Reprinted with permission of Elsevier.

My central interest is not with whether global shape representation poses an insuperable challenge for deep neural nets (although see Bowers et al., 2023). Rather, I suggest that studies aiming to evaluate whether newly sighted subjects perceive global shape, or only collections of local features, may take inspiration from this parallel line of research. Would newly sighted subjects exposed to a series of normal squares and normal circles immediately recognize the images in figure 3 as more “square-like” or “circle-like” in a manner resembling normally sighted participants? Or do they respond more like AlexNet? If the latter, then this may support Schwenkler’s local feature-matching hypothesis.

In any case, the point I wish to emphasize is that in stark contrast to Issue 1, studies of this sort *may* generate evidence bearing on Issue 2. Suppose that newly sighted subjects were to *succeed* in such global shape matching tasks while still failing Held et al.’s T-V task. Then, a natural explanation of this combination of results would be that newly sighted subjects successfully represent global shape in both vision and touch, but these representations are rationally disconnected.

Notice that this would not be the *only* possible explanation of failure in the T-V task. Perhaps newly sighted subjects’ visual and haptic representations of shape *are* rationally connected, but the

subjects fail to appreciate the identity of seen and felt shape for some other reason. As noted earlier, they might be distracted by extraneous differences between their overall visual and haptic experiences (such as the appearance of color in vision or thermal features in touch), and unable to maintain focus on the shared geometrical content between the two. However, such results would at least supply evidence that it is possible to perceive a single shape property through sight and touch while coherently doubting that this is the case. The evidence would be defeasible, inasmuch as further findings might favor a rival hypothesis, but almost all psychological evidence is defeasible in *this* sense. Molyneux's test may be *helpful* in settling Issue 2 even if it is not an *experimentum crucis*.

Here, then, is the crucial contrast between Issue 1 and Issue 2. Because the Match Principle is false, studies of newly sighted subjects (no matter how well-designed) are unlikely to be informative about Issue 1. Conversely, studies of newly sighted subjects *may* be informative about Issue 2. For the latter to be the case, however, we require stronger evidence than we currently possess that newly sighted subjects represent *shape* in vision, rather than mere collections of local contour features. Recent work on the visual capacities of deep neural nets offers a helpful parallel in investigating this issue.

4. Radical externalism and rational connection

Let's grant that newly sighted subjects are potentially informative about Issue 2—whether it is impossible for a fully reflective subject who both sees and feels the same shape property to coherently doubt that this is the case. I now consider this issue's broader significance within philosophy of perception, and specifically whether it marks a fault line in debates about the metaphysics of perceptual experience. John Campbell (1996) argues that prominent “radically externalist” theories of perceptual experience are committed to resolving Issue 2 in the affirmative. After laying out Campbell's argument in section 4.1, I argue in section 4.2 that the argument fails because radical externalism entails no such

commitment. More generally, while Issue 2 is independently interesting, it cannot be used to adjudicate debates about the fundamental nature of perceptual experience.

4.1. Radical externalism and Campbell's argument

Campbell (1996) characterizes “radical externalism” as the view that “what makes one’s consciousness of shape is the fact that one is using a neural system whose role is to pick up the shape properties of the objects in one’s environment” (p. 302). In particular, according to radical externalism, the phenomenal character of spatial experience is determined by the mind-independent spatial properties presented by the experience. Moreover, the properties so presented are those that are “picked up” or (perhaps better) encoded by the neural systems recruited when undergoing the experience.

Campbell argues that if radical externalism is correct, then the phenomenal character of shape perception must be the same in vision and touch. Thus:

Insofar as we are externalist about shape perception, we have to think of it as amodal. For insofar as we are externalist about shape perception, we have to think of experience of shape as a single phenomenon, in whatever sense-modality it occurs, individuated by the external geometrical property. For it is in fact the very same properties that are being perceived by sight as by touch. (...) For the radical externalist...there is no difference in the phenomenal character of shape experience in sight and touch. (pp. 303-304)¹⁹

He further claims that if shape phenomenology is the same between vision and touch, then there should be no room for rational doubt about whether visual and haptic experiences of *square* present the same worldly property:

[I]nsofar as we are externalist about the character of shape perception, then there is nothing in the character of the experience itself to ground a doubt as to whether it is the same

¹⁹ In Campbell’s terminology, to say that shape perception is “amodal” is to say that it has a common sensory phenomenal character in both vision and touch. For alternative ways of understanding the notion of amodality, see Spence and Di Stefano (2024). In particular, note that Campbell is *not* attributing to radical externalists the controversial claim that shape perception requires deploying abstract concepts lacking in sensory character, but he is attributing *more* than the uncontroversial claim that vision and touch converge on a common external property (viz., shape) (see Spence & Di Stefano, 2024, p. 13).

properties that are being perceived through vision as through touch. (...) The sameness of property perceived in sight and touch is transparent to the subject, and cross-modal transfer is a rational phenomenon. (pp. 303-304)

In other words, radical externalists are committed to a positive answer to Issue 2. We can reconstruct Campbell's reasoning as follows:

1. If radical externalism is correct, then visual and haptic experiences are alike in shape phenomenology.
2. If two perceptual experiences are alike in shape phenomenology, then (assuming human-like reasoning capacities) there is no room for rational doubt about whether they present the same shape properties.
3. Therefore, if radical externalism is correct, then there is no room for rational doubt about whether visual and haptic experiences present the same shape properties.

Note that if this argument is sound, then visual and haptic presentations of the same shape property must be rationally connected not just in normally sighted human perceivers, but also in any other perceiver for whom radical externalism is true. Presumably, if radical externalism is true of normally sighted perceivers, then it is also true of newly sighted perceivers, and of conscious perceivers more generally. Accordingly, Campbell's argument establishes a connection between newly sighted perceivers and the metaphysics of perceptual experience. If a fully reflective newly sighted perceiver can coherently doubt whether their visual and haptic experiences present the same shape properties, then radical externalism is false.

While I will be concentrating on premise 1 in the next subsection, I should flag that premise 2 in Campbell's argument is open to doubt. For, in the above passage, Campbell transitions without argument from the claim that visual and haptic shape experiences share phenomenal character to the conclusion that "the sameness of property perceived in sight and touch is transparent to the subject." It is unclear why this transition is legitimate. For, even if radical externalism is *in fact* true of perceptual

experience, the subject may not *know* that it is true. Accordingly, a fully reflective subject might coherently wonder whether two experiences with the same spatial phenomenal character really present the same spatial properties. In any event, I believe that the argument fails even if we grant premise 2, so I won't press this concern in what follows.

A wide family of theories is committed to radical externalism as Campbell understands it, since many accept that the phenomenal character of perceptual experience is fully determined by the worldly objects and properties presented to one in having the experience. But different family members construe the notion of presentation differently.

One view construes presentation as a species of *representation* (Byrne, 2001; Dretske, 1995; Tye, 1995, 2000). Thus, according to what Pautz (2021, p. 142) labels *response-independent representationalism*, the phenomenal character of an experience just is its property of representing a certain array of mind-independent objects and properties. Given response-independent representationalism (hereafter simply “representationalism”), any difference in phenomenal character between two experiences requires a difference in the properties or objects represented by the experiences. This inspires a *prima facie* compelling line of thought mirroring Campbell's argument: Assume that representationalism is true. Then, if sight and touch represent the very same shape properties, then at least as regards the phenomenology of *shape*, their phenomenal characters should be identical. And if the two experiences are indeed identical in shape phenomenology, then (assuming premise 2) any fully reflective subject should be in position to determine beyond any rational doubt that they present the same shape properties.

Alternatively, naïve realists analyze experiential presentation in terms of a non-representational relation of *acquaintance* (Brewer, 2011; Fish, 2009; French & Phillips, 2020; Martin, 2004). The relation of acquaintance is non-representational insofar as it does not admit of perceptual error (Byrne & Green, 2023). One cannot be acquainted with an object's being red if it is not red, while one *can*

experientially *represent* an object as red when it is not red. However, setting aside cases of illusion, naïve realism can be developed in ways that closely resemble representationalism. For example, Fish (2009) holds that the phenomenal character of an experience just is its property of “acquainting the subject with a selection of the facts that inhabit the tract of the environment the subject perceives” (p. 75). On this view, any difference in phenomenal character must be accompanied by some difference in the mind-independent facts with which one is acquainted.

Fish construes the facts presented in experience as “object-property couples” (2009, p. 52), in which a property is instantiated by an object or a relation is instantiated by multiple objects. Thus: “[W]hen we see an object (a) by seeing one of its properties (its F-ness), we can say that we see a particular fact—the fact of a’s being F” (Fish, 2009, p. 52). Now consider the visual and haptic experience of a square object, o. The following line of thought may seem *prima facie* compelling. Both the visual experience and the haptic experience acquaint the subject with the fact of o’s being square. Assuming that the phenomenal character of experience just is its property of acquainting the subject with such mind-independent facts, then at least as regards their shape phenomenology, the two experiences should be phenomenally identical. And if the experiences are identical in shape phenomenology, then any fully reflective subject should be in position to determine beyond any rational doubt that they present the same shape properties.

Thus, various representationalists and naïve realists endorse radical externalism.²⁰ Moreover, one can give a *prima facie* compelling argument that on either of these views, Issue 2 should be resolved

²⁰ Specifically, radical externalism is entailed by forms of representationalism and naïve realism according to which any two experiences that present the same worldly objects or properties also share the same phenomenology. Note, however, that not all representationalists or naïve realists endorse this claim. For example, Fregean representationalists permit the *mode of presentation* of an object or property to play a role in determining phenomenal character (Schellenberg, 2018, ch. 4; Thompson, 2010). Such theorists may hold that vision and touch both represent shape, but under different modes of presentation, producing cross-modal differences in phenomenal character. And on the naïve realist side, several have argued that phenomenal character is determined not only by *which* mind-independent elements we are perceptually acquainted with, but also by the *ways* we are acquainted with them (Beck, 2019; French & Phillips, 2020; Logue, 2012; Martin, 1998, p. 175; Sethi, in press). Such theorists might similarly claim that vision and touch both present shape, but in different ways, producing differences in phenomenal character.

affirmatively. Thus, if Campbell is right, then Issue 2 can be used to adjudicate a central debate about the nature of perceptual experience—viz., whether the character of experience is determined “externally” by the objects and properties we perceive, or “internally” by neuro-functional properties of our brains.

4.2. Why radical externalism is not committed to rational connection


I now argue that we should reject the foregoing argument. Radical externalism does not entail that visual and haptic states are exactly alike in shape phenomenology, so Campbell’s argument fails at the first step. Thus, irrespective of its independent interest, Issue 2 cannot be used to adjudicate the debates about whether phenomenal character is determined externally or internally. Furthermore, radically externalist theories of experience are compatible with any outcome of Molyneux’s test.

While I am focusing on Campbell’s argument, the idea that differences in spatial phenomenology across modalities might pose a problem for radical externalism is not unique to Campbell (see also Block, 1996; Lopes, 2000; O’Dea, 2006). For, various authors have argued that because we perceive some of the same spatial properties through multiple modalities, radically externalist theories entail that the phenomenology associated with these properties should also be shared across modalities. So, if it isn’t (as many deem plausible), then radical externalism is false. If I am right, however, then radical externalism has no such consequence. Thus, because Campbell’s argument fails at the first step, its failure infects any argument that takes cross-modal differences in shape phenomenology to challenge radical externalism.

My main points are these: First, the human perception of shape is *compositional* insofar as our perception (or perceptual representation) of global shape properties is built from our perception (or perceptual representation) of simpler shape primitives (Biederman, 1987; Green, 2019, 2023; Hafri et

al., 2023; Hummel, 2000, 2013; Lande, 2024). Second, there is no reason to think that vision and touch necessarily employ the same class of shape primitives. Third, radical externalism is compatible with the compositional character of shape perception. Fourth, if radical externalism is correct, then differences in the shape primitives employed within vision and touch should be expected to ground differences in the phenomenal character of shape within the two modalities, and such differences may also block any rational connection between visual and haptic shape perception. Thus, radical externalism is compatible with differences in shape phenomenology in vision and touch, and also with a negative resolution to Issue 2.

One sign that human shape perception is compositional is that we can easily appreciate the similarities between objects composed of the same parts in different spatial arrangements (Arguin & Saumier, 2004; Behrmann et al., 2006; Hummel, 2000), and visual priming also transfers between such objects (Cacciamani et al., 2014). These data can be explained on the assumption that some of the same capacities are exercised when we perceive distinct objects composed of common parts. For instance, when we perceive the shape in figure 4a, we employ some of the same perceptual capacities that we employ when we perceive the shape in figure 4b. Given a representational analysis of these capacities (e.g., Schellenberg, 2018), a natural hypothesis is that the visual system produces representations of the mid-sized parts of each object (the oval, rectangle, and trapezoid), and that these representations are repeated for the two figures, although they are composed with different representations of spatial relations, accounting for the two shapes' discriminability.



Figures 4a (left) and 4b (right). Two shapes composed of the same parts in different spatial arrangements.

Nonetheless, while mundane observations such as these indicate *that* shape perception is compositional, they leave open the *way* in which it is compositional. Specifically, they leave open which primitives are employed by systems of perceptual shape representation, and how these primitives are combined to yield representations of global shape. Here, contemporary theories of shape representation offer a range of alternatives (Green, 2019, 2023; Lande, 2024). Different schemes of shape description offer competing ways of encoding the same global shape properties.

One salient distinction among shape representation schemes concerns the dimensionality of their primitives—specifically, whether they describe shape in terms of the arrangement of one-dimensional features like edges or vertices (Ullman & Basri, 1991), two-dimensional features like planar surface shapes (Leek et al., 2009), or three-dimensional features like volumetric cones or cylinders (Biederman, 1987; Marr & Nishihara, 1978). One complication in comparing schemes that differ in the dimensionality of their primitives is that such schemes often differ in *which* shape properties they represent (e.g., 2D surface shapes versus 3D volumes), and not merely in the *ways* these properties are represented. Thus, I will restrict my focus to competing schemes for representing

two-dimensional (surface or image) shape. Such schemes function to encode the same class of shape properties, but via different primitives and modes of combination.

One crucial distinction is between “contour-based” and “region-based” schemes (Zhang & Lu, 2004). Contour-based schemes encode shape via an ordered sequence of representations of contour points or segments, where representations adjacent in the sequence function to encode spatially adjacent contour segments. One such scheme is the *centroid distance function*, which encodes a shape’s contour points in terms of their distances from the shape’s centroid (figure 5; Zhang & Lu, 2004). Here, the global representation is an ordered sequence of representations of distances from the centroid to contour points. Baker et al.’s (2021) constant-curvature approach offers another contour-based scheme in which an extended, closed contour is represented by an ordered sequence of representations of constant curvature segments. Each such segment is described by its turning angle (a measure of curvature) and length.

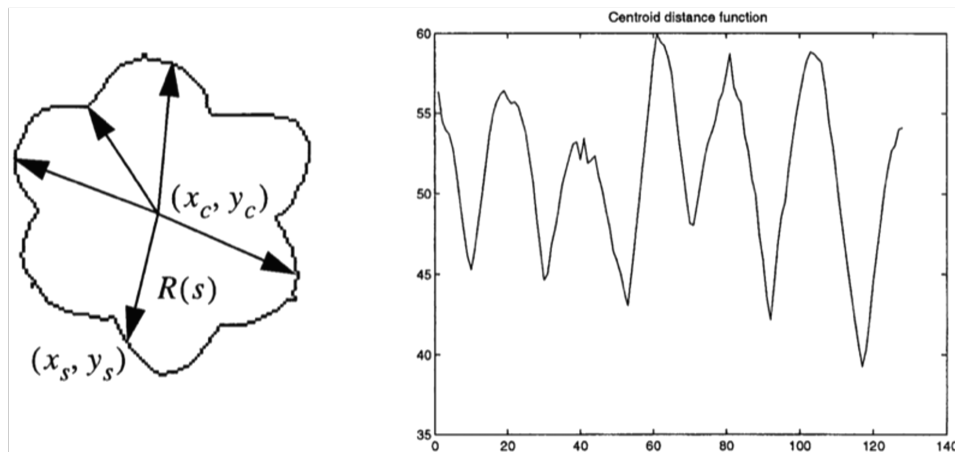


Figure 5. Centroid distance function (right) of the 2D shape on the left, composed of a sequence of representations of the distance from the centroid to points on the contour of the shape, where representations adjacent in the sequence encode spatially adjacent contour points. *Source:* Ma & Manjunath (1999). Reprinted with permission of Springer Nature.

Region-based schemes encode shape via properties of global shape regions, such as their medial axes or axes of symmetry. A 2D shape's medial axis is composed of the set of points in the interior of the shape that have two or more nearest neighbors on the shape's boundary (Blum, 1973). When plotted, the medial axis resembles a skeleton, and different branches of the skeleton correspond roughly to different parts of the shape (see figure 6).²¹ One prominent elaboration of the medial-axis scheme is Feldman and Singh's (2006) skeletal theory. On this theory, 2D shape is encoded via a hierarchical tree in which each "node" of the tree corresponds to a separate part of the shape, and parts are segmented in accordance with axis branches. The tree is organized into a "root" node encoding the main body of the shape, followed by descendants encoding parts that protrude from the main body (e.g., "limbs").




Figure 6. Medial axes of the human hand. Notice that separate axis branches tend to correspond to separate parts of the hand (e.g., fingers). *Source:* Feldman & Singh (2006). Copyright (2006) National Academy of Sciences.

These alternative schemes exhibit differences both in which geometrical properties they take as primitive (i.e., the basic "building blocks" from which global shape properties are described), and their mode of combination (e.g., an ordered sequence versus a hierarchical tree). Thus, while the constant curvature and skeletal approaches take turning angles (say) as a primitive feature, the centroid distance function need not. More broadly, these differences illustrate that the same shape property

²¹ Unfortunately, the raw medial axis is highly sensitive to contour noise, sprouting spurious branches that do not correspond to perceptually significant parts of the shape. In response to this difficulty, theorists have offered strategies for "smoothing" the medial axis to remove such spurious branches. For instance, Feldman and Singh (2006) impose a prior on skeletons that assigns higher probability to smoother axes and penalizes extra axis branches.

can be encoded by means of different descriptive contents, depending on the scheme of shape representation employed.

In light of this, radical externalists are not committed to the claim that any two experiences that present the same shape property must share the same shape phenomenology. Consider first the representationalist who holds that the phenomenal character of an experience just is its property of representing a certain content involving an array of worldly objects and properties. We've seen that different arrays of worldly properties can suffice to specify the same global shape (e.g., centroid distances versus constant curvature segments). Accordingly, the radical externalist can permit two experiences that represent the same global shape to differ in shape phenomenology, provided that they differ in the primitive properties through which they compositionally encode that global shape. Accordingly, since shape representations in distinct modalities might differ in their geometrical primitives, it is compatible with representationalism that visual and haptic experiences of the same shape property may differ in shape phenomenology.

Similar remarks hold for naïve realism. Recall Fish's view that the phenomenal character of a perceptual experience is its property of acquainting the subject with a selection of mind-independent facts. It seems plausible that, on this view, perceptual acquaintance with complex structural facts (e.g., an object's being human-shaped) might be constituted by acquaintance with more primitive facts (e.g., its being composed of a head, two arms, two legs, and a torso). If so, it is open to naïve realists of Fish's stripe to hold that two experiences might both acquaint the subject with an object's being square, but by means of acquainting the subject with different primitive facts about that object (e.g., its being composed of a set of contour points exhibiting a particular centroid distance function, or its being composed of four straight segments of equal length, meeting at four right angles). Such differences in the primitive facts with which we are acquainted may ground differences in shape phenomenology between vision and touch.

Thus, we should reject premise 1 in the above argument. It is not true that if radical externalism is correct, then visual and haptic experiences must be alike in shape phenomenology. I suggest that the apparent plausibility of this premise derives from insufficient reflection on the compositional character of shape perception.

However, one might still be tempted to think that if radical externalism is correct, then there should necessarily be a rational connection between visual and haptic experiences of the same shape property. For, even if a given shape property is “described” differently in vision and touch, it should nonetheless be possible for a fully reflective subject to *figure out* that the two descriptions are necessarily coextensive. And if so, there should be no room for rational doubt regarding whether the two experiences present the same property.

In fact, toward the end of his paper Campbell concedes that differences in how shape properties are described in vision and touch might yield minor differences in the phenomenology of shape, even if radical externalism is true. However, he insists that such differences still would not block a rational connection between the experiences:

[I]t will be possible for different geometrical descriptions to be given of the very same shapes in sight than in touch; indeed, two different visual perceptions of the same shape may give different geometrical descriptions of it, as when one object is a rotated version of another, similarly shaped thing. In this case it may still be informative to be told that the shapes are the same; so if vision and touch give different geometrical descriptions of the same shape, it may still be informative to be told that it is the same shape one is seeing as touching. But given the unity of the underlying, externally constituted geometry of the two senses, it will be possible for the perceiver to determine *a priori* that it is the same shape that is in question. (Campbell, 1996, p. 317)

Campbell doesn't tell us what he means by the “unity of the underlying, externally constituted geometry of the two senses,” or why radical externalism entails the presence of such a unified geometry. (Indeed, it is an open question whether there is even a unified geometry characterizing the representation of spatial properties within vision; see Wagner, 2006). In any case, however, I see no good reason to accept that identities between shape properties perceptually represented by means of

distinct representation schemes must be appreciable *a priori*, if this means that they must be appreciable by reasoning alone. Or, at least, I see no reason to accept it unless we also adopt a controversial view of the relation between perception and thought.

Campbell seems to think that if two perceptual states offer distinct but necessarily coextensive descriptions of a given shape property, then it should be possible for a rational subject to “reason with” these descriptions to recover their equivalence. Such reasoning might, for instance, begin with a representation of the centroid distance function and derive, step-by-step, an equivalent skeletal representation. However, doing so would presumably require deploying *concepts* of the properties encoded by the two descriptions (e.g., CENTROID, TURNING ANGLE, AXIS POINT, BOUNDARY POINT, etc.). Yet it is consistent with radical externalism that perceptual states are non-conceptual—i.e., that a perceptual state might represent or acquaint one with a given property even though one lacks the concept of that property (Block, 2023; Heck, 2000; Tye, 1995, 2000). So, even if a subject is fully reflective, there is no guarantee that she will have the conceptual repertoire needed to determine the equivalence of distinct visual and haptic descriptions of a given shape property.²²

Moreover, those who hold that perceptual representation is at least partly non-conceptual typically pair this position with a view about the formats or “codes” employed in perception. Specifically, perception employs at least some codes that are qualitatively different from those employed in reasoning and inference, and that are only eligible for restricted forms of computation (Block, 2023; Carey, 2009; Quilty-Dunn, 2020). Accordingly, if certain perceptual shape descriptions are couched in a format that is ineligible to participate directly in discursive geometrical reasoning,

²² In fact, even if perception *is* wholly conceptual (in a certain sense), it *still* wouldn’t follow that the coextension of distinct perceptual representations of a given shape property must be transparent to the subject. Suppose that a subject must possess concepts of every property treated as primitive within any scheme of shape representation her perceptual system employs (e.g., TURNING ANGLE, AXIS POINT, etc.). Still, the subject may be unable to derive equivalences between two shape descriptions unless she *also* possesses certain *linking* concepts necessary for relating the primitives in one scheme to those in the other. But even if perception is wholly conceptual, there is no guarantee that the subject will possess such linking concepts.

then regardless of the concepts a subject happens to possess, she may not be able to reason with the two perceptual descriptions to work out their equivalence.²³ So, in addition to conceptual limitations, constraints imposed by representational format or cognitive architecture may also prevent a fully reflective subject from deriving equivalences between distinct visual and haptic shape descriptions. Crucially, however, all this is compatible with the view that the phenomenal character of visual and haptic experiences is fully determined by the external properties presented by those experiences.

I conclude that radical externalism about perceptual experience is compatible with differences in shape phenomenology across modalities, and also with a negative resolution to Issue 2: Even if the phenomenology of perception is wholly determined by the mind-independent objects and properties we perceive, it would not follow that any rational subject must be able to derive the equivalence of shape properties presented in different modalities. Thus, Issue 2 cannot be used to adjudicate debates about the fundamental nature of perceptual experience.

5. Conclusion

In considering Molyneux's question, it is important to distinguish the question itself, which poses a concrete experimental test of a newly sighted perceiver, from various theoretical issues to which the question's answer might be relevant. One broad topic concerns whether there is a "rational connection" between the representations of shape deployed in sight and touch—that is, a connection that would be transparent to any rational subject given the proper time, attention, and motivation.

I've suggested that we should distinguish two versions of the rational connection question. The first version, which I've called Issue 1, concerns whether the visual and haptic states through

²³ Note that even if the subject has this inability, she may still be able to use the perceptual shape descriptions for *some* psychological tasks—e.g., imagery tasks like mental rotation or scanning. But such tasks would presumably involve altering a shape's encoding *within* the relevant descriptive scheme (e.g., its represented orientation within that scheme). They would not require "translating" between distinct schemes of shape description.

which *normally sighted perceivers* apprehend shape, and which mediate our capacities for cross-modal recognition, are rationally connected. While this issue has significant implications for understanding the architecture of our sensory modalities and the forms of interaction they exhibit, I've argued that newly sighted subjects are largely irrelevant to it. Thus, Molyneux's question (at least interpreted literally) does not bear interestingly on the relation between visual and haptic shape representations in normally sighted perceivers.

The second version, which I've called Issue 2, concerns whether visual and haptic presentations of shape are *necessarily* rationally connected—i.e., whether it is *impossible* for a fully reflective perceiver to perceive the same shape property through both sight and touch while coherently doubting that this is so. I've argued that newly sighted subjects may bear on Issue 2, but that whether they do depends on whether they visually represent shape *at all*—and this is indeed an open question.

Nevertheless, irrespective of its independent interest, I've argued that the outcome of Issue 2 cannot settle debates about the metaphysics of perceptual experience. For even if the character of perceptual experience is fully determined by the worldly objects and properties we perceive, two senses may present or “describe” the same global shape property using different geometrical primitives, and equivalences between these modality-specific descriptions may not be derivable by reasoning alone.

Where does this leave the question that Molyneux initially posed to Locke? I suggest that while the question does retain some theoretical interest, it is far less significant for understanding our perception and cognition of shape than one might have thought. The question's answer cannot tell us much about the nature of cross-modal recognition (or, indeed, the nature of shape concepts) in normally sighted individuals, and neither can it be used to adjudicate between rival views of the nature of perceptual experience.²⁴

²⁴ For helpful comments and discussion, I am grateful to Ian Phillips, Kevin Lande, and two anonymous reviewers for this journal.

References

- Alary, F., Duquette, M., Goldstein, R., Chapman, C. E., Voss, P., La Buissonnière-Ariza, V., & Lepore, F. (2009). Tactile acuity in the blind: a closer look reveals superiority over the sighted in some but not all cutaneous tasks. *Neuropsychologia*, *47*(10), 2037-2043.
- Alary, F., Goldstein, R., Duquette, M., Chapman, C. E., Voss, P., & Lepore, F. (2008). Tactile acuity in the blind: a psychophysical study using a two-dimensional angle discrimination task. *Experimental Brain Research*, *187*, 587-594.
- Amedi, A., Malach, R., Hendler, T., Peled, S., & Zohary, E. (2001). Visuo-haptic object-related activation in the ventral visual pathway. *Nature Neuroscience*, *4*(3), 324-330.
- Amedi, A., Raz, N., Azulay, H., Malach, R., & Zohary, E. (2010). Cortical activity during tactile exploration of objects in blind and sighted humans. *Restorative Neurology and Neuroscience*, *28*(2), 143-156.
- Arguin, M., & Saumier, D. (2004). Independent processing of parts and of their spatial organization in complex visual objects. *Psychological Science*, *15*(9), 629-633.
- Baker, N., & Elder, J. H. (2022). Deep learning models fail to capture the configural nature of human shape perception. *iScience*, 104913.
- Baker, N., Garrigan, P., & Kellman, P. J. (2021). Constant curvature segments as building blocks of 2D shape representation. *Journal of Experimental Psychology: General*, *150*(8), 1556-1580.
- Baker, N., Lu, H., Erlikhman, G., & Kellman, P. J. (2018). Deep convolutional networks do not classify based on global object shape. *PLoS Computational Biology*, *14*(12), e1006613.
- Baker, N., Lu, H., Erlikhman, G., & Kellman, P. J. (2020). Local features and global shape information in object classification by deep convolutional neural networks. *Vision Research*, *172*, 46-61.
- Beck, O. (2019). Rethinking naive realism. *Philosophical Studies*, *176*, 607-33.
- Behrmann, M., Peterson, M.A., Moscovitch, M., & Suzuki, S. (2006). Independent representation of parts and the relations between them: Evidence from integrative agnosia. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(5), 1169-1184.
- Berkeley, G. (1709/1965). *Essay Towards a New Theory of Vision*. In D. M. Armstrong (ed.), *Berkeley's Philosophical Writings*. New York: Collier Books.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*(2), 115-117.
- Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human perception and performance*, *19*(6), 1162-1182.
- Block, N. (1996). Mental paint and mental latex. *Philosophical Issues*, *7*, 19-49.
- Block, N. (2023). *The Border Between Seeing and Thinking*. Oxford: Oxford University Press.
- Blum, H. (1973). Biological shape and visual science (Part I). *Journal of Theoretical Biology*, *38*(2), 205-287.
- Bowers, J. S., Malhotra, G., Dujmović, M., Montero, M. L., Tsvetkov, C., Biscione, V., ... & Blything, R. (2023). Deep problems with neural network models of human vision. *Behavioral and Brain Sciences*, *46*, e385.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.
- Bruno, M., & Mandelbaum, E. (2010). Locke's answer to Molyneux's thought experiment. *History of Philosophy Quarterly*, *27*(2), 165-180.
- Byrne, A. (2001). Intentionalism defended. *The Philosophical Review*, *110*(2), 199-240.
- Byrne, A., & Green, E. J. (2023). Whither naïve realism?—I. *Philosophical Perspectives*, *37*(1), 49-68.
- Cacciamani, L., Ayars, A.A., & Peterson, M.A. (2014). Spatially rearranged object parts can facilitate perception of intact whole objects. *Frontiers in Psychology*, *5*, 1-11.

- Campbell, J. (1996). Molyneux's question. *Philosophical Issues*, 7, 301-318.
- Campbell, J. (2005). Information processing, phenomenal consciousness, and Molyneux's question. In J. L. Bermúdez (ed.), *Thought, Reference, and Experience: Themes from the Philosophy of Gareth Evans*, pp. 195-219. Oxford: Oxford University Press.
- Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.
- Chaisilprungraung, T., German, J., & McCloskey, M. (2019). How are object shape axes defined? Evidence from mirror-image confusions. *Journal of Experimental Psychology: Human Perception and Performance*, 45(1), 111-124.
- Cheng, T. (2015). Obstacles to testing Molyneux's question empirically. *i-Perception*, 6(4), 1-5.
- Cheselden, W. (1728). An account of some observations made by a young gentleman, who was born blind, or lost his sight so early, that he had no remembrance of ever having seen, and was couch'd between 13 and 14 years of age. *Philosophical Transactions of the Royal Society of London*, 35(402), 447-450.
- Connolly, K. (2013). How to test Molyneux's question empirically. *i-Perception*, 4(8), 508-510.
- Copenhaver, R. (2014). Berkeley on the language of nature and the objects of vision. *Res Philosophica*, 91(1), 29-46.
- Degenaar, M. (1996). *Molyneux's Problem: Three Centuries of Discussion on the Perception of Forms*. Dordrecht: Kluwer.
- Desmarais, G., Meade, M., Wells, T., & Nadeau, M. (2017). Visuo-haptic integration in object identification using novel objects. *Attention, Perception, & Psychophysics*, 79, 2478-2498.
- Di Stefano, N., & Spence, C. (2023). Perceptual similarity: Insights from crossmodal correspondences. *Review of Philosophy and Psychology*. DOI: 10.1007/s13164-023-00692-y.
- Dretske, F. I. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Drewes, J., Goren, G., Zhu, W., & Elder, J. H. (2016). Recurrent processing in the formation of shape percepts. *Journal of Neuroscience*, 36(1), 185-192.
- Elder, J. H. (2018). Shape from contour: Computation and representation. *Annual Review of Vision Science*, 4, 423-450.
- Erdogan, G., Chen, Q., Garcea, F. E., Mahon, B. Z., & Jacobs, R. A. (2016). Multisensory part-based representations of objects in human lateral occipital cortex. *Journal of Cognitive Neuroscience*, 28(6), 869-881.
- Erdogan, G., & Jacobs, R. A. (2017). Visual shape perception as Bayesian inference of 3D object-centered shape representations. *Psychological Review*, 124(6), 740-761.
- Erdogan, G., Yildirim, I., & Jacobs, R. A. (2015). From sensory signals to modality-independent conceptual representations: A probabilistic language of thought approach. *PLoS Computational Biology*, 11(11), e1004610.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429-433.
- Evans, G. (1985). Molyneux's question. In A. Phillips (ed.), *Gareth Evans: Collected Papers*. Oxford: Clarendon Press. Reprinted in A. Noë & E. Thompson (eds.), *Vision and Mind*. (2002). Cambridge, MA: MIT Press, pp. 319-349.
- Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences*, 103(47), 18014-18019.
- Fine, I., Wade, A. R., Brewer, A. A., May, M. G., Goodman, D. F., Boynton, G. M., Wandell, B. A., & MacLeod, D. I. (2003). Long-term deprivation affects visual perception and cortex. *Nature Neuroscience*, 6(9), 915-916.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- French, C., & Phillips, I. (2020). Austerity and illusion. *Philosophers' Imprint*, 20, 1-19.
- Gallagher, S. (2005). *How the Body Shapes the Mind*. Oxford: Oxford University Press.

- Glenney, B. R. (2013). Philosophical problems, cluster concepts, and the many lives of Molyneux's question. *Biology & Philosophy*, 28, 541-558.
- Gori, M. (2015). Multisensory integration and calibration in children and adults with and without sensory and motor disabilities. *Multisensory Research*, 28(1-2), 71-99.
- Gori, M., Del Viva, M., Sandini, G., & Burr, D. C. (2008). Young children do not integrate visual and haptic form information. *Current Biology*, 18(9), 694-698.
- Gori, M., Giuliana, L., Sandini, G., & Burr, D. (2012). Visual size perception and haptic calibration during development. *Developmental Science*, 15(6), 854-862.
- Gori, M., Sandini, G., Martinoli, C., & Burr, D. (2010). Poor haptic orientation discrimination in nonsighted children may reflect disruption of cross-sensory calibration. *Current Biology*, 20(3), 223-225.
- Green, E. J. (2019). On the perception of structure. *Noûs*, 53(3), 564-592.
- Green, E. J. (2022a). Representing shape in sight and touch. *Mind & Language*, 37(4), 694-714.
- Green, E. J. (2022b). The puzzle of cross-modal shape experience. *Noûs*, 56(4), 867-896.
- Green, E. J. (2023). A pluralist perspective on shape constancy. *The British Journal for the Philosophy of Science*. DOI: 10.1086/727427.
- Gregory, R. L., & Wallace, J. G. (1963). Recovery from early blindness: A case study. *Experimental Psychology Society Monograph*, 2, 65-129.
- Hafri, A., Green, E. J., & Firestone, C. (2023). Compositionality in visual perception. *Behavioral and Brain Sciences*, 46, e277.
- Heck, R. (2000). Nonconceptual content and the "space of reasons". *The Philosophical Review*, 109(4), 483-523.
- Heimler, B., Behor, T., Dehaene, S., Izard, V., & Amedi, A. (2021). Core knowledge of geometry can develop independently of visual experience. *Cognition*, 212, 104716.
- Held, R., Ostrovsky, Y., de Gelder, B., Gandhi, T., Ganesh, S., Mathur, U., & Sinha, P. (2011). The newly sighted fail to match seen with felt. *Nature Neuroscience*, 14(5), 551-553.
- Hopkins, R. (2005). Molyneux's question. *Canadian Journal of Philosophy*, 35(3), 441-464.
- Huber, E., Webster, J. M., Brewer, A. A., MacLeod, D. I., Wandell, B. A., Boynton, G. M., Wade, A. R., & Fine, I. (2015). A lack of experience-dependent plasticity after more than a decade of recovered sight. *Psychological Science*, 26(4), 393-401.
- Hummel, J. E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. In E. Dietrich & A. Markman (eds.), *Cognitive Dynamics: Conceptual Change in Humans and Machines*, pp. 157-185. Hillsdale, NJ: Lawrence Erlbaum.
- Hummel, J. E. (2013). Object recognition. In D. Reisburg (ed.), *Oxford Handbook of Cognitive Psychology*, pp. 32-46. Oxford: Oxford University Press.
- Humphreys, G. W., & Quinlan, P. T. (1988). Priming effects between two-dimensional shapes. *Journal of Experimental Psychology: Human Perception and Performance*, 14(2), 203-220.
- James, T. W., Humphrey, G. K., Gati, J. S., Menon, R. S., & Goodale, M. A. (2002). Differential effects of viewpoint on object-driven activation in dorsal and ventral streams. *Neuron*, 35(4), 793-801.
- Kalenine, S., Pinet, L., & Gentaz, E. (2011). The visual and visuo-haptic exploration of geometrical shapes increases their recognition in preschoolers. *International Journal of Behavioral Development*, 35(1), 18-26.
- Kellman, P. J., & Fuchser, V. (2023). Visual completion and intermediate representations in object formation. In A. Mroczko-Wasowicz & R. Grush (eds.), *Sensory Individuals: Unimodal and Multimodal Perspectives*, pp. 55-76. Oxford: Oxford University Press.
- Kellman, P. J., Guttman, S. E., & Wickens, T. D. (2001). Geometric and neural models of object

- perception. In T. F. Shipley & P. J. Kellman (eds.), *From Fragments to Objects: Segmentation and Grouping in Vision*, pp. 183-246. Amsterdam: Elsevier.
- Lacey, S., Hall, J., & Sathian, K. (2010). Are surface properties integrated into visuohaptic object representations? *European Journal of Neuroscience*, *31*(10), 1882-1888.
- Lacey, S., Peters, A., & Sathian, K. (2007). Cross-modal object recognition is viewpoint-independent. *PLoS One*, *2*(9), e890.
- Lacey, S., & Sathian, K. (2014). Visuo-haptic multisensory object recognition, categorization, and representation. *Frontiers in Psychology*, *5*, 730.
- Lande, K. J. (2024). Compositionality in perception: A framework. *WIREs Cognitive Science*.
- Lawson, R. (2009). A comparison of the effects of depth rotation on visual and haptic three-dimensional object recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(4), 911-930.
- Leek, E. C., Reppa, I., Rodriguez, E., & Arguin, M. (2009). Surface but not volumetric part structure mediates three-dimensional shape representation: Evidence from part-whole priming. *Quarterly Journal of Experimental Psychology*, *62*(4), 814-830.
- Lescroart, M. D., & Biederman, I. (2013). Cortical representation of medial axis structure. *Cerebral Cortex*, *23*(3), 629-637.
- Levin, J. (2008). Molyneux's question and the individuation of perceptual concepts. *Philosophical Studies*, *139*(1), 1-28.
- Levin, J. (2018). Molyneux's question and the amodality of spatial experience. *Inquiry*, *61*(5-6), 590-610.
- Locke, J. (1694/1975). *An Essay Concerning Human Understanding*. P. H. Nidditch (ed.). Oxford: Oxford University Press.
- Logue, H. (2012). Why naive realism? *Proceedings of the Aristotelian Society*, *112*, 211-37.
- Lopes, D. M. M. (2000). What is it like to see with your ears? The representational theory of mind. *Philosophy and Phenomenological Research*, *60*(2), 439-453.
- Ma, W. Y., & Manjunath, B. S. (1999). Netra: A toolbox for navigating large image databases. *Multimedia Systems*, *7*, 184-198.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London, B: Biological Sciences*, *200*, 269-294.
- Martin, M. G. F. (1998). Setting things before the mind. In A. O'Hear (ed.), *Current Issues in Philosophy of Mind*. Cambridge: Cambridge University Press.
- Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, *120*, 37-89.
- Masson, H. L., Bulthé, J., Op de Beeck, H. P., & Wallraven, C. (2016). Visual and haptic shape processing in the human brain: unisensory processing, multisensory convergence, and top-down influences. *Cerebral Cortex*, *26*(8), 3402-3412.
- Mathes, B., & Fahle, M. (2007). Closure facilitates contour integration. *Vision Research*, *47*(6), 818-827.
- Matthen, M., & Cohen, J. (2020). Many Molyneux questions. *Australasian Journal of Philosophy*, *98*(1), 47-63.
- May, E., Arach, P., Kishiki, E., Geneau, R., Maehara, G., Sukhai, M., & Hamm, L. M. (2022). Learning to see after early and extended blindness: A scoping review. *Frontiers in Psychology*, *13*, 954328.
- McKyton, A., Ben-Zion, I., Doron, R., & Zohary, E. (2015). The limits of shape recognition following late emergence from blindness. *Current Biology*, *25*(18), 2373-2378.
- Newell, F. N., Ernst, M. O., Tjan, B. S., & Bühlhoff, H. H. (2001). Viewpoint dependence in visual and haptic object recognition. *Psychological Science*, *12*(1), 37-42.

- Noppeney, U., Friston, K. J., Ashburner, J., Frackowiak, R., & Price, C. J. (2005). Early visual deprivation induces structural plasticity in gray and white matter. *Current Biology*, *15*(13), R488-R490.
- Norman, J. F., Norman, H. F., Clayton, A. M., Lianekhammy, J., & Zielke, G. (2004). The visual and haptic perception of natural object shape. *Perception & Psychophysics*, *66*(2), 342-351.
- O’Callaghan, C. (2019). *A Multisensory Philosophy of Perception*. Oxford: Oxford University Press.
- Ocelli, V., Lacey, S., Stephens, C., John, T., & Sathian, K. (2016). Haptic object recognition is view-independent in early blind but not sighted people. *Perception*, *45*(3), 337-345.
- O’Dea, J. W. (2006). Representationalism, supervenience, and the cross-modal problem. *Philosophical Studies*, *130*(2), 285-295.
- Orlov, T., Raveh, M., McKyton, A., Ben-Zion, I., & Zohary, E. (2021). Learning to perceive shape from temporal integration following late emergence from blindness. *Current Biology*, *31*(14), 3162-3167.
- Ostrovsky, Y., Meyers, E., Ganesh, S., Mathur, U., & Sinha, P. (2009). Visual parsing after recovery from blindness. *Psychological Science*, *20*(12), 1484-1491.
- Park, C., & Kim, S. (2014). Haptic perception accuracy depending on self-produced movement. *Journal of Sports Sciences*, *32*(10), 974-985.
- Pautz, A. (2021). *Perception*. London: Routledge.
- Peissig, J. J., & Tarr, M. J. (2007). Visual object recognition: Do we know more now than we did 20 years ago? *Annual Review of Psychology*, *58*, 75-96.
- Poggio, T., & Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, *343*(6255), 263-266.
- Postma, A., Zuidhoek, S., Noordzij, M. L., & Kappers, A. M. (2008). Haptic orientation perception benefits from visual experience: Evidence from early-blind, late-blind, and sighted people. *Perception & Psychophysics*, *70*(7), 1197-1206.
- Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press.
- Putzar, L., Hötting, K., Rösler, F., & Röder, B. (2007). The development of visual feature binding processes after visual deprivation in early infancy. *Vision Research*, *47*(20), 2616-2626.
- Quilty-Dunn, J. (2020). Perceptual pluralism. *Noûs*, *54*(4), 807-838.
- Quinlan, P. T., & Humphreys, G. W. (1993). Perceptual frames of reference and two-dimensional shape recognition: Further examination of internal axes. *Perception*, *22*(11), 1343-1364.
- Reid, T. (1764/1997). *An Inquiry into the Human Mind*. Edinburgh: Edinburgh University Press.
- Sacks, O. (1995). *An Anthropologist on Mars*. New York, NY: Vintage Books.
- Sann, C., & Streri, A. (2007). Perception of object shape and texture in human newborns: Evidence from cross-modal transfer tasks. *Developmental Science*, *10*(3), 399-410.
- Schellenberg, S. (2018). *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Schwenkler, J. (2012). On the matching of seen and felt shape by newly sighted subjects. *i-Perception*, *3*(3), 186-188.
- Schwenkler, J. (2013). Do things look the way they feel? *Analysis*, *73*(1), 86-96.
- Schwenkler, J. (2019). Molyneux’s question within and across the senses. In T. Cheng, O. Deroy, & C. Spence (eds.), *Spatial Senses: Philosophy of Perception in an Age of Science*. London: Routledge.
- Sethi, U. (forthcoming). A (qualified) defense of diaphaneity. In O. Beck & F. Masrour (eds.), *The Relational View of Perception: New Essays*. London: Routledge.
- Sourav, S., Bottari, D., Kekunnaya, R., & Roder, B. (2018). Evidence of a retinotopic organization of early visual cortex but impaired extrastriate processing in sight recovery individuals. *Journal of Vision*, *18*(3), 22.

- Spelke, E. (2022). *What Babies Know*. Oxford: Oxford University Press.
- Spence, C., & Di Stefano, N. (2024). What, if anything, can be considered an amodal sensory dimension? *Psychonomic Bulletin & Review*. DOI: 10.3758/s13423-023-02447-3.
- Streri, A., & Gentaz, E. (2003). Cross-modal recognition of shape from hand to eyes in human newborns. *Somatosensory & Motor research*, 20(1), 13-18.
- Tarr, M. J., & Pinker, S. (1990). When does human object recognition use a viewer-centered reference frame? *Psychological Science*, 1(4), 253-256.
- Theurel, A., Frileux, S., Hatwell, Y., & Gentaz, E. (2012). The haptic recognition of geometrical shapes in congenitally blind and blindfolded adolescents: is there a haptic prototype effect?. *PLoS one*, 7(6), e40251.
- Thompson, B. (2010). The spatial content of experience. *Philosophy and Phenomenological Research*, 81(1), 146-184.
- Tye, M. (1997). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Ueda, Y., & Saiki, J. (2012). Characteristics of eye movements in 3-D object learning: comparison between within-modal and cross-modal object recognition. *Perception*, 41(11), 1289-1298.
- Ullman, S. & Basri, R. (1991). Recognition by linear combinations of models. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13, 992-1006.
- van Cleve, J. (2007). Reid's answer to Molyneux's question. *The Monist*, 90(2), 251-270.
- Wagner, M. (2006). *The Geometries of Visual Space*. Mahwah, NJ: Lawrence Erlbaum.
- Wiesel, T. N., & Hubel, D. H. (1963). Single-cell responses in striate cortex of kittens deprived of vision in one eye. *Journal of Neurophysiology*, 26(6), 1003-1017.
- Woodman, G. F. (2010). A brief introduction to the use of event-related potentials in studies of perception and attention. *Attention, Perception, & Psychophysics*, 72(8), 2031-2046.
- Yildirim, I., & Jacobs, R. A. (2013). Transfer of object category knowledge across visual and haptic modalities: Experimental and computational studies. *Cognition*, 126(2), 135-148.
- Zhang, D., & Lu, G. (2004). Review of shape representation and description techniques. *Pattern Recognition*, 37(1), 1-19.