Empirical Explanations of the Laws of Appearance[*]

Penultimate draft (forthcoming, *Journal of Philosophy*)

Abstract: It is widely thought that there are limits to how things can perceptually appear to us. For instance, nothing can appear both square and circular, or both pure red and pure blue. Adam Pautz has dubbed such constraints "laws of appearance." But if the laws of appearance obtain, then what explains them? Here I examine the prospects for an empirical explanation of the laws of appearance. First, I challenge extant empirical explanations that appeal purely to the format of perceptual representation. I then develop a hybrid approach, on which the laws are explained not merely by format, but by two further factors: ecological constraints imposed by our environments, and computational constraints embodied by our perceptual systems. While the hybrid approach implies that the laws of appearance are contingent, I argue that this implication is empirically defensible, since even some of the most intuitively compelling laws have real-world counterexamples.

Are there limits to how things can appear to us in perceptual experience? Intuitively, yes. No surface can appear both pure blue and pure red, or both square and circular. Such intuitions are widespread. The idea that a single thing might appear two distinct colors or shapes is commonly taken to be absurd, and often dismissed without argument. Thus, C. D. Broad briefly considers the hypothesis that a region might be "sensibly pervaded by some shade of red, and…at the same time sensibly pervaded by some shade of green," and swiftly remarks: "This, I suppose, would be admitted to be impossible" (1925: 164).

Recently, Adam Pautz has called attention to various constraints of this sort, dubbing them *laws of appearance* (Pautz 2017, 2020, 2021). These laws include the following:

> *Exclusion law.* An individual cannot experientially represent that a single surface has two distinct pure colors, such as pure red and pure green. Likewise, an individual cannot experientially represent that the same object has distinct shapes, such as round and square.
>
> *Berkeley's law.* (1) An individual cannot experientially represent that something has a color without also experientially representing that it takes up space in some way. (2) Conversely, an individual cannot experientially represent that something takes up space in some way (e.g., being circular) without also experientially representing a qualitative difference (e.g., a *white* circle on a *black* background).

[*] Thanks to Alex Byrne and Adam Pautz for valuable comments and discussion of earlier drafts of this paper, and to Matt Duncan, Gabe Greenberg, Kevin Lande, Fiona Macpherson, Adam Pautz, and Jake Quilty-Dunn for helpful conversations about this material. Thanks also to an audience that the 2024 World Congress of Philosophy in Rome. Finally, I am indebted to two anonymous reviewers for detailed comments that greatly improved the paper.

> *Perspectival law.* An individual cannot experientially represent merely *that there is a cube somewhere in reality*, without any "perspectival content" about its location and apparent shape from "here." (Pautz 2021, 131)

Each of these laws embodies a regularity wherein experiences of one sort of property either require or preclude experiences of other sorts of properties. For example, the Exclusion Law states that experiencing a surface as having one color precludes experiencing it as having a distinct color. And the Perspectival Law states (*inter alia*) that experiencing an object as having some shape requires experiencing it as having some egocentric location (location "from here").

Though less broadly accepted than the Exclusion Law, Berkeley's Law and the Perspectival Law also have enduring appeal. As its name suggests, the former derives from Berkeley's view that we cannot form ideas of primary qualities like shape that fully abstract away from secondary qualities like color or texture (Berkeley 1906: 47). And as regards the Perspectival Law, the claim that the perceptual representation of shape requires representation of egocentric location (or other egocentric spatial properties) has been endorsed by Gareth Evans among others (Evans 1985/2002; Peacocke 1992: 61-67; Burge 2014).

Assuming that the laws of appearance hold, how are they to be explained? Pautz argues that representationalists about perceptual experience have no easy answer this question, and he rejects several possible explanatory strategies. While I agree that the laws call out for explanation, I am less pessimistic about the prospects for giving one. This paper critiques a recent strategy for explaining the laws, then outlines an alternative approach.

While all the laws merit discussion, I'll focus primarily on the Exclusion Law, which is the most widely endorsed of the laws, and where Pautz places greatest emphasis. For convenience, I divide this law into the *Color* and *Shape* Exclusion Laws. The former states that we cannot visually experientially represent (i.e., have a visual experience that represents) one surface as having two distinct colors at the same level of determinacy (e.g., pure red and pure green). The latter states that we cannot visually experientially represent one object as having two distinct shapes at the same level of determinacy (e.g., spherical and cubical).[1] Both laws are coherence principles mandating the consistency of experiences of a single dimension. And both are commonly accepted. Even those who think that a surface can *be* both red and green typically grant that no surface can *appear* both red and green simultaneously, at least to human perceivers (Kalderon 2007: 572). And it is often treated as a basic desideratum on theories of the perspectival character of perception that they avoid commitment to pervasive violations of the Shape Exclusion Law (Siewert 2006: 4-5; Green & Schellenberg 2018: 4).

The Exclusion Laws plausibly have *some* modal force. We do not ordinarily see things as both circular and square, and it seems unlikely that this is a mere accidental generalization akin to "all the coins in my pocket are dimes." However, there is disagreement about whether the laws are *metaphysically necessary* truths about perception or *contingent laws* of human psychology, which might be violated in non-human perceivers or in humans under atypical circumstances. Following Pautz, I'll call these views *Necessitism* and *Contingentism*, respectively.

Pautz takes the laws of appearance to challenge *representationalism*: the view that phenomenal properties just are properties of experientially representing certain contents. For, regardless of whether the laws are necessary or contingent, we would like an explanation of them, but representationalism doesn't obviously provide one. Representationalism is *existence-neutral*: Because an experience might represent the presence of an *F* without there being any *F* in your surroundings, it might appear to you that something is *F*, even if nothing is *F* (Pautz 2020, 257). Thus, if representationalism is true, then we cannot infer from the fact that nothing could *be* both red and

[1] We might also formulate generalizations of both Exclusion Laws requiring harmony across levels of determinacy. Specifically, if an experience represents an object as having some shape or color property *F*, then any other shape or color property it attributes to the object must be either a determinate or a determinable of *F*. Thus, if an experience represents a surface as scarlet, then it may also represent it as red, but not as green. I suspect that anyone who accepts the Color or Shape Exclusion Laws would also accept these generalized variants.

green that nothing could *appear* this way. Moreover, representationalists construe perceptual experience as a species of representational mental state, but such laws do not seem to constrain other representational mental states. After all, someone might *suppose* or mistakenly *believe* that some surface is both red and green. Thus, representationalism does not a priori entail that experience must obey the laws of appearance.

Here, representationalism is disadvantaged relative to sense-datum theory. Sense-datum theorists claim that when it appears to you that something is *F*, you are aware of a sense-datum that really is *F*. Accordingly, if nothing (not even a sense-datum) could *be* both red and green, then nothing could *appear* this way (Pautz 2020, 259). This tempting explanation of the Exclusion Laws is unavailable to representationalists, or to anyone else who embraces existence-neutrality. Thus, beyond their intrinsic interest, the laws of appearance bear on fundamental debates about the metaphysics of perceptual experience.

This paper considers *empirical explanations* of the laws of appearance. By an "empirical explanation," I mean an explanation that appeals to empirically discoverable facts about our perceptual systems or the environments we inhabit. Conversely, *a priori explanations* purport to explain the laws via a priori knowable facts about perception. For example, someone might argue that we can establish sense-datum theory a priori, and deduce the laws therefrom. However, the laws are unlikely to receive a priori explanations if representationalism is correct. Because representationalists need a viable empirical explanation of the laws, and my sympathies lie broadly with representationalism (Green 2016; Byrne & Green 2023), I will focus on empirical explanations.

Empirical explanations are compatible with either Necessitism or Contingentism, depending on the modal status of the facts in the explanation. Suppose we explain a law of appearance by appeal to some empirically discovered feature *F* of human perception. If *F* is an *essential* feature of perception (that is, *F* is instantiated by the perceptual system of any possible perceiver), then that law would be metaphysically necessary. In discovering that perception is *F*, we would be unearthing a necessary a posteriori truth about perception. If *F* is a *contingent* feature of perception (that is, there are possible perceivers whose perceptual systems lack *F*), then the law might turn out to be contingent as well. Still, certain empirical explanations fit better with Contingentism than Necessitism, since it may just be highly implausible that the facts in the explanation constitute essential features of perception.

Extant empirical explanations of the laws of appearance have appealed solely to the *format* of perceptual representation, specifically its alleged iconic or depictive format. My first aim is to clarify why iconic format might be thought to play this explanatory role. Then I argue that pure format-based explanations rely on an implicit but unwarranted *uniqueness assumption* about the representational underpinnings of perceptual experience. Next, I identify two further sources of evidence that may be employed in explaining the laws of appearance: *ecological constraints* imposed by the environments we occupy, and format-extrinsic *computational constraints* on the flow of information within our perceptual systems. A viable empirical explanation should, I contend, enlist all three forms of evidence. I call this the *hybrid approach* to explaining the laws of appearance. I then argue that the hybrid approach leads to Contingentism. However, I contend that we should accept this consequence, since even principles as intuitively compelling as the Shape Exclusion Law have real-world counterexamples.

1. Format-based Explanations

Format-based explanations of the laws of appearance have recently been floated by Tye (2020), Sainsbury (2023), and Morgan (2023). According to Sainsbury, a format-based explanation of some law of appearance "would show the law to rest on format features of the underlying vehicles" (2023, 2937). I will interpret the format-based approach as claiming that the laws of appearance rest *solely* on format properties of the representational vehicles underlying perceptual experience. I'm unsure whether anyone would endorse a position quite this strong, but it is an independently interesting view worth considering, particularly since format properties have dominated discussions of empirical explanations of the laws, and it is the primary type of empirical explanation that Pautz himself considers (2021: 132-133).

The *format* of a representational system consists in principles governing the relationship between the syntactic properties and the contents of representations belonging to that system. Syntactic properties can be understood as nonsemantic features of representational vehicles to which computational operations on those vehicles are sensitive, and which determine a representation's well-formedness within the system. Syntactic properties include a representation's decomposition into constituents alongside their ordering or arrangement. Often, systems of representation require well-formed representations to contain certain types of constituents in certain arrangements. Thus, "Bob ate pizza" is well-formed in English, while "Ate Bob pizza" is not. I leave open whether syntactic properties should be identified with functional or physiological properties of mental representations.

Relations between syntactic properties and contents capture the way that a system of representation "codes" for its contents. Consider the contrast between Arabic and binary numeral systems. Both systems represent number. However, the Arabic representation of the number 53 ("53") and its binary representation ("110101") have different syntactic properties, including different degrees of structural complexity. Moreover, constituents of the representations map to different contents: Constituents of an Arabic numeral encode the number's decomposition into powers of 10, while constituents of a binary numeral encode its decomposition into powers of 2 (Marr 1982: 20). Thus, Arabic and binary systems encode number in different formats.

Broad categories of formats (e.g., discursive, iconic, analog) are characterized by generalizations over the relations between the syntactic properties and the contents of representations within each category (Haugeland 1991). Thus, on a popular view, a system displays *analog* format just in case there is a monotonic mapping from some syntactic magnitude (M1) of the representational vehicles within the system to a magnitude (M2) that they represent, such that when a vehicle's value along M1 increases, the value it represents along M2 systematically increases or systematically decreases (Maley 2011; Beck 2019; Peacocke 2019). Accordingly, mercury thermometers display analog format because they exhibit a syntactic magnitude (mercury height) that varies monotonically with the magnitude they represent (temperature).

1.1. Against Tye-Sainsbury Format-Based Explanations

How might format explain the laws of appearance? Several authors have suggested that the laws result from some broadly *depictive* property that perceptual representations allegedly exhibit. For instance, Sainsbury (2023) notes that parallels of some of Pautz's laws hold for pictures (no picture can depict a single region as both pure green and pure blue), and suggests that format-based explanations might capitalize on an analogy to pictures, though he acknowledges that perceptual representations are not literally pictorial (2937). Block (2023) suggests that certain of Pautz's laws (namely, the No Logical Structure Law) may result from the analog format of perception, though he claims that others "are not truths at all" (199). Finally, Tye (2020, 67-68) proposes that the "array" format of perceptual representation explains several of the laws. I'll focus on Tye's proposal, since it is the most developed.

Tye characterizes array format as follows:

> A…plausible view is that visual experiences have the structure of arrays, as understood in computer science, where these arrays are made up of cells that are dedicated to lines of sight in the field of view and that contain symbols (that is, simple, primitive representations) themselves dedicated to representing local features

> on those lines of sight, for example, the color of a tiny surface patch lying there, its distance away and whether there is part of an edge on it. On this view, visual experiences have a matrix-like structure… (2020, 68; see also Tye 2000, 71-72)

Thus, arrays are composed of cells, with adjacent cells dedicated to adjacent lines of sight. Each cell contains symbols encoding the distance of any surface visible along the relevant line of sight alongside "qualitative" information about the color or texture of the surface. The shape of an object is represented in terms of the directions, distances, and orientations of its visible surfaces. For example, a slanted circle and head-on oval might fall along the same lines of sight, and thus be assigned the same cells of an array, but would differ in the distance or orientation properties attributed to them by these cells.

Regarding the laws of appearance, Tye writes:

> The hypothesized matrix structure for visual experiences, under suitable further elaboration, can also be used to explain the laws of appearance. For example, the reason why nothing can look to be both pure blue and bluish green is that for each cell, there is only a single symbol dedicated to color and if the symbol represents one color on a given line of sight, it cannot simultaneously represent another there. Similarly, something cannot look both cubical and spherical because that would require inconsistencies in how the symbols in cells represent edges; and so on. (2020, 68)

Tye argues that discursive format offers no parallel explanation of the laws of appearance, and that this favors the view that visual experience is array-like over the view that it is discursive.

However, there is a critical defect in Tye's explanation. According to Tye, if experience has an array structure, then we should expect the Exclusion Laws to hold. We can't represent a single region as both pure blue and bluish green because, says Tye, "for each cell, there is only a single symbol dedicated to color" (68). The problem is that this restriction on color attribution does not follow from the preceding characterization of array format. According to Tye, an array is a structure composed of cells dedicated to lines of sight, where each cell contains "symbols…themselves dedicated to representing local features" (68). But if this is all that array format requires, then a single cell could easily contain symbols for incompatible color properties, violating the Exclusion Law.

Since each cell represents a complex conjunction of features (e.g., glossy, blue and sloping) in a given direction, the mere fact that each cell is assigned to just one direction does not entail that it cannot represent an incompatible conjunction, such as blue, green, and glossy (see figure 1). The requirement that each array cell contain only one symbol for color is an additional constraint beyond array format, as Tye characterizes it. Similar remarks apply to the Shape Exclusion Law. To prevent the attribution of incompatible shape properties to a single surface (e.g., circular and oval), we would need to add that each cell holds only one symbol for distance or orientation. Again, this constraint does not derive from Tye's characterization of array format. It is a brute limit on the storage capacity of array cells.

2 m away Glossy Pure blue Bluish green	Edge here	3 m away Matte Pure red
2 m away Glossy Bluish green		3 m away Matte Pure red

Figure 1. An array-like representation violating the Color Exclusion Law (top left cell).

Thus, because Tye's explanation of the Exclusion Laws depends crucially on format-extrinsic constraints on the storage capacity of array cells, he has not shown how the Exclusion Laws derive from the *format* of arrays. It is easy to construct arrays meeting Tye's specifications that violate them.

1.2. A Modified Format Explanation

Nonetheless, using Kulvicki's (2015) notion of *incompatibility classes*, we can modify Tye's notion of array format to ensure that array representations *cannot* violate the Exclusion Laws. As Kulvicki observes, map-like representations encode the locations of objects or properties by placing syntactic markers—"locatable features"—at corresponding locations on a map. For instance, a map of the earth might place the locatable feature *blueness* at all map regions corresponding to regions where the earth contains water, and *greenness* at all map regions corresponding to regions where the earth contains land. Blueness and greenness are incompatible syntactic properties: A single map region cannot instantiate both of them. Moreover, the properties that these syntactic properties encode—land and water—are also incompatible. A properly designed map obeys what Kulvicki calls the *incompatibility constraint*: Incompatibility classes of locatable features serve to represent classes of properties that are also mutually incompatible.[2]

Thus, beyond claiming that visual arrays contain cells dedicated to particular regions or directions, the format-based theorist should add that they obey the incompatibility constraint. Specifically: (i) each cell of a visual array is assigned to some direction or region of the visual field, (ii) no two cells are assigned to the same direction or region, and (iii) the syntactic properties instantiated by a cell are mutually incompatible if and only if the properties they represent are mutually incompatible.[3] Thus, because *colors themselves* are mutually incompatible, the syntactic properties used to represent them must also be mutually incompatible. Similar remarks hold for the representation of distance and orientation, leading to the Shape Exclusion Law.

[2] See Kulvicki (2015, 155). Maps violating this constraint can be constructed, but they are defective in their expressive power.

[3] Note that some might find condition (iii) dubious, since it excludes the possibility of *any* mismatches between the syntactic properties of an array and the values of the dimensions they pick out (e.g., cases where incompatible syntactic properties just happen to pick out the same value). Of course, the condition might be weakened in various ways to accommodate exceptions.

Notice that the incompatibility constraint does not require an array to *reproduce* the properties it represents, just as maps of Earth needn't literally contain water and land. It requires only that mutually incompatible classes of syntactic properties correspond to perceptible dimensions with mutually incompatible values. Thus, maps obeying the incompatibility constraint arguably supply a better analogy than pictures for proponents of the format-based approach.

A proponent of this account might argue that the Exclusion Laws result purely from the format of perceptual representation: Because visual experiences are underpinned by (and inherit their contents from) visual arrays, and the format of a visual array prevents it from representing two incompatible colors in the same region, visual experience cannot represent this type of scenario either. Call this the *modified format explanation*. Unfortunately, while this line of thought is tempting, I contend that it is flawed.

2. Against the Modified Format Explanation

The problem is that while the incompatibility constraint can explain why an *individual* visual array cannot represent incompatible colors in the same region, we cannot infer that visual experience is similarly constrained unless we also assume that visual experience inherits its content solely from an individual array. For the modified format explanation to succeed, we must assume that some array *uniquely* subserves visual experience. However, this *uniqueness assumption* is dubious.

If a map of the Earth obeys the incompatibility constraint, then it cannot represent a single region as containing both land and water. However, nothing prevents such conflicts from arising *between* maps. If one cartographer knows about an island in the Pacific of which a second is ignorant, then they may produce maps that conflict regarding the presence of land in the relevant region. Likewise, if the visual system produces multiple arrays each encoding color across the visual field, then even if the incompatibility constraint holds within each array, nothing prevents violations of the

Color Exclusion Law across arrays. Moreover, if the visual system produces both iconic, array-like representations and non-iconic, discursive representations (Quilty-Dunn et al. 2023), then there is nothing to prevent conflicts between iconic and non-iconic representations regarding the color or shape of a given surface.

Pautz (2021) raises a related concern for pure format-based approaches:

> When you view a moving blue thing, your…neural representations of color and movement, which might be in separate brain areas, are "bound" together, where this is some kind of functional-computational relation. Why couldn't there be possible experiencers whose subpersonal representations of distinct colors…could be "bound together" in this way, just as our subpersonal representations of color and motion can be bound together? (133)

I suggest that advocates of the format-based approach can offer a *partial* answer to Pautz's binding question. If the visual system produces arrays that satisfy the incompatibility constraint, then the incompatibility of syntactic properties encoding distinct colors or spatial properties may explain why an individual array cannot violate the Exclusion Laws. Conversely, the syntactic properties encoding colors and motion trajectories are compatible, so these features are freely co-attributable. Nevertheless, format alone cannot explain why violations of the Exclusion Laws do not occur *across* arrays, or between arrays and non-arrays.

While the uniqueness assumption *could* be correct, I see little reason to grant it without a convincing argument in its favor. Moreover, there are strong reasons to doubt it (see also Martinez & Nanay 2024).

One reason derives from the complexity of shape experience. Visual arrays of the sort described earlier provide a purely *viewer-centered* representation of shape, in which an object's shape is encoded by means of the viewer-centered directions, distances, and orientations of its visible surfaces. Accordingly, array representations of shape are volatile across changes in perspective, which drastically alter these properties (Bennett 2012; Todd & Petrov 2022). Conversely, *object-centered* representations are based on an object's intrinsic axes, such as medial axes or axes of symmetry, and remain approximately stable through various perspective shifts (Erdogan & Jacobs 2017; Chaisilprungraung et al. 2019).

There is evidence that visual shape representation is not entirely array-like. First, the visual system seems to recover abstract or qualitative shape properties (for example, affine or topological properties) which are not explicitly encoded by standard array structures like Marr's 2.5-D sketch but are explicitly encoded by alternative representational schemes (Bennett 2012; Green 2017, 2023a; Todd & Petrov 2022). There is also growing evidence that the visual system generates object-centered shape representations based on an object's *skeletal* or *medial axis structure*, which resemble the skeletal or "stick-figure" structure of the object (Feldman & Singh 2006). Such representations encode the layout of an object's boundaries via their spatial relations to its skeletal axes. Skeletal axes are relied on when discriminating or categorizing shapes across changes in orientation and surface features (Lowet et al. 2018). Moreover, canonical visual brain areas code for skeletal structure in a manner that is roughly invariant to changes in viewpoint (Lescroart & Biederman 2013).

Green (2019, 2022) argues that while viewer-centered, array-like representations plausibly *contribute* to our perceptual experience of shape (see also Briscoe 2008; Lande 2018), the *overall* experience is best explained by a combination of viewer-centered and object-centered representations. Viewer-centered representations underlie aspects of shape phenomenology that change with shifts in perspective, while object-centered representations underlie aspects that remain more stable. Moreover, object-centered representations may be structured more like a rooted tree or directed graph than an array (Feldman & Singh 2006).

Viewer-centered shape representations encode certain properties that object-centered shape representations do not, such as viewer-centered depth and orientation. However, this doesn't prevent each from denoting the same shape properties. Rather, I suggest that they simply pick out these properties in different ways, using different primitives or modes of combination. Viewer-centered representations describe shape properties via primitives for viewer-centered distance and direction, while object-centered representations describe these same properties via primitives for intrinsic axis structure and spatial relations to points along the axis. By analogy, consider the descriptions "Closed figure with three angles" and "Closed figure with three sides." Each encodes a feature that the other does not, but both are ways of specifying the same shape. I suggest that viewer-centered and object-centered representations bear a parallel relation. Both encode shape properties, but specify them via different primitive features.

The compresence of viewer-centered and object-centered shape representations in visual phenomenology is exemplified by cases of shape ambiguity. Consider the isosceles triangle and the square-diamond in figure 2. 45°-rotations differentially affect one's experiences of these figures, producing a bigger change to the square's appearance than the triangle's. The triangle appears different at different orientations, but the square appears *more* different. Why? Plausibly, because shape experience is jointly underpinned by viewer-centered and object-centered representations. Viewer-centered representations ground aspects of shape appearance that vary in both cases, while object-centered representations ground aspects that vary in the square case but not the triangle case. Specifically, when an object has multiple symmetry axes available for constructing an object-centered reference frame, the visual system may favor the axis closest to vertical in viewer-centered coordinates. Since the square has multiple symmetry axes while the triangle doesn't, the visual system generates different object-centered representations of the square at different orientations, but the same object-centered representation of the triangle. Consistent with this, there is evidence that isosceles triangles visually prime their 45°-rotated counterparts, while squares do not (Humphreys & Quinlan 1988).

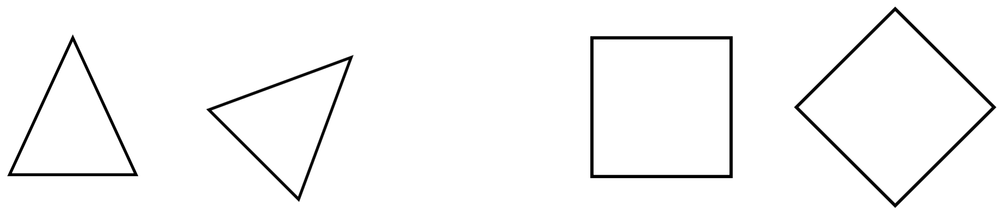

Figure 2. An isosceles triangle and a square/diamond.

Thus, contra the uniqueness assumption, visual shape experience is likely *not* uniquely underpinned by a single viewer-centered array, but rather by multiple representational structures that code for shape in different ways (for instance, using distinct reference frames and geometrical primitives). This poses a problem for the modified format explanation. While format constraints may prevent violations of the Shape Exclusion Law *within* viewer-centered arrays, they do not obviously prevent violations of the law *between* viewer-centered and object-centered representations (for example, representing an object as circular within a viewer-centered frame but as elliptical within an object-centered frame). In fact, I'll argue in section 4 that subtle violations of this sort do occur.[4]

I do not claim that viewer-centered and object-centered representations float entirely free of one another in the visual system. First, there may be computational interactions between them. Information encoded by one representation may be accessed by computations that produce the other.[5] Second, the visual system may represent spatial relations between select elements encoded by

[4] While certain object-centered representations encode geometrical properties at higher levels of abstraction, such representations can also encode determinate or "metric" shape properties (Green 2023a). (Roughly, metric shape properties are those that are preserved only under rigid transformations or uniform scaling—see Lee et al. (2012).) Here, I am primarily concerned with the potential for conflict between viewer-centered and object-centered representations of metric shape, and these are the sorts of conflicts uncovered by Li (2009), discussed below. Nonetheless, even if object-centered and viewer-centered representations do encode shape at strictly different levels of abstraction, the representations could generate violations of the generalized versions of the Exclusion Laws requiring harmony both within and across levels of abstraction (see note 1).

[5] Marr (1982) famously proposed that viewer-centered representations of depth and orientation serve as the primary inputs to computations of object-centered representations. However, this view has since been challenged (Bennett 2012; Todd & Petrov 2022, 15).

the two representations without transforming those representations into a common format. For example, according to the *coordinate-system orientation representation theory* (COR) (McCloskey et al. 2006; Gregory & McCloskey 2010), object-centered representations are linked to viewer-centered representations by encoding correspondences between the axes of the frames together with the angular separation between them. For example, suppose that viewer-centered representations incorporate perpendicular axes aligned with the viewer's up-down and left-right directions. Then, when you view the tilted triangle in figure 2, perception might represent (i) a correspondence between the triangle's symmetry axis and your own vertical axis, (ii) the direction in which the former axis is rotationally offset from the latter, and (iii) the magnitude of the angle between them (Gregory & McCloskey 2010: 125).

If the COR model is correct, then we should expect some degree of unity between viewer-centered and object-centered aspects of shape phenomenology, despite their heterogeneous representational bases. The COR scheme binds an object's object-centered shape representation with a representation of its viewer-centered position and orientation, explaining why these aspects of shape appearance seem to "go together" and to characterize a single individual.[6]

However, while the COR scheme establishes a link between object-centered and viewer-centered representations, the degree of integration is minimal because there is only one point of direct contact between them: Perception simply represents how one representation's axes are oriented with respect to the other's. Despite this, the representations may continue to encode shape in radically different formats (e.g., using different primitives or composition rules), and inconsistencies between the two representations may be allowed to persist, even if inconsistencies

[6] While the COR scheme relates object-centered and viewer-centered representations within vision, it may fail to relate certain shape representations across modalities. This could explain why an object's viewer-centered and object-centered visual appearances seem more unified than (say) its viewer-centered visual appearance and hand-centered haptic appearance. However, certain object-centered aspects of shape phenomenology may be tightly unified across modalities (Green 2022).

within them are prohibited by format. Accordingly, while the COR model establishes a form of *weak integration* between viewer-centered and object-centered representations, this doesn't suffice to rescue the modified format explanation.

The fundamental challenge for the modified format explanation is that format-based constraints on feature binding that apply within representational structures may fail to apply across structures, particularly when these structures code for the relevant properties in fundamentally different ways. And there is strong reason to believe that perceptual experience is underpinned by multiple representational structures. Accordingly, format constraints on individual perceptual representations do not immediately carry over to experience.

One might object that the foregoing remarks apply only to shape, not color. However, the evidence suggests that color experience also depends on a heterogeneous variety of representational structures. Color vision seems to comprise separate systems of "color-for-coloring" and "color-for-form" (Akins & Hahn 2014; Davies 2021). The former recovers monadic color qualities, while the latter recovers the locations and qualities of edges between surfaces, plausibly generating representations of color relations like *redder than* (Davies 2021). These systems dissociate in cerebral achromatopsia, which involves loss of color-for-coloring but partially spared color-for-form (Kentridge et al. 2004). Thus, it is doubtful that color experience is uniquely subserved by a single visual array. If so, syntactic constraints on color attribution within individual arrays cannot explain the Color Exclusion Law.[7]

[7] One might suggest that because the Exclusion Law only concerns monadic color experience, a pure format-based explanation is viable as long as the "color-for-coloring" system generates a single array. However, this approach would still fail to explain other coherence principles that seem just as compelling as the Color Exclusion Law. Suppose that color relations and monadic colors are encoded by separate representations at different stages of processing. Then format principles alone may permit apparent monadic colors and color relations to conflict (e.g., experiencing *A* as pure blue, *B* as pure green, but *B* as bluer than *A*).

One might suggest that while *early* vision encodes features across distinct representational structures, all of them eventually feed into a "master icon" meeting the incompatibility constraint, and perceptual experience is uniquely subserved by this master icon. Proponents of this model might take inspiration from Treisman's feature integration theory (FIT), on which individual features (e.g., redness, blueness, verticality) are coded in separate, uncoordinated maps in early vision, which eventually feed into a "master map" of locations (Treisman 1988). Proponents of the modified format explanation might suggest that FIT's master map meets the incompatibility constraint and is the sole basis of perceptual experience, guaranteeing compliance with the Exclusion Laws.

However, this proposal requires both that FIT is correct, *and* that perceptual experience depends solely on the master map. But FIT faces difficulties. Specifically, the claims that feature binding occurs at a single processing stage, and that binding requires attention, are empirically dubious (Humphreys 2016; Quilty-Dunn 2023). Moreover, even if FIT is true, the second claim is likely false. FIT's master map does not represent particular features at unattended locations, but only degrees of overall salience (Treisman 1988, 203). Assuming that experience represents *some* particular features outside of one's attentional focus, FIT's master map cannot uniquely subserve experience. Moreover, if visual experience depends partially on FIT's earlier maps, the problem for the modified format explanation resurfaces. Nothing in the *format* of the redness and blueness maps prevents them from attributing redness and blueness, respectively, to a single region. Thus, FIT cannot motivate the modified format explanation.

3. The Hybrid Approach

3.1. Beyond Format: Ecological and Computational Constraints

I contend that a viable empirical explanation of the laws of appearance should appeal not merely to the format of perceptual representation, but to two further sources of evidence: *ecological constraints* imposed by our environments, and format-extrinsic *computational constraints* on how perceptual representations are formed and integrated. According to the *hybrid approach*, the laws of appearance are explained by all three factors operating in tandem.

In defending the hybrid approach, I am arguing that a *general form* of explanation is correct. Given present knowledge, we cannot be sure exactly which explanation of this form is correct. Moreover, the relative explanatory significance of the three factors will surely vary from one law of appearance to another. These are matters for further investigation. Furthermore, despite rejecting *pure* format-based explanations of the Exclusion Laws, I do not recommend that we *ignore* format properties (see sections 3.2-3.3). However, we should abandon any ambition of *reading off* the laws of appearance from the format of perceptual representation in the way that parallel laws can be derived from the format of pictorial representation.

The fundamental problem for the modified format explanation is that the same dimension is often represented via heterogeneous representational structures each of which contributes to the contents of experience. Even if format constraints ensure compliance with the Exclusion Laws within each representation, they may fail to ensure compliance across representations. By analogy, even if individual maps are format-bound to obey the incompatibility constraint, pairs of maps generally are not. Fortunately, both ecological and computational constraints can fill this gap.

First, ecological constraints. Often, when perception represents a single property via separate representations, those representations are produced through distinct channels that capitalize on different cues to the property. For example, vision estimates slant from both stereoscopic and textural cues (Hillis et al. 2002). Here, two visual estimates of a dimension are formed through separate channels. However, the *environment* helps to secure their agreement. If principles of optics, together with lawlike generalizations about the environment (e.g., surfaces tend to be homogeneously textured), ensure that two cues are both highly reliable indicators of values along a distal dimension, then estimates derived from these cues will tend to converge in normal conditions.

The production of viewer-centered and object-centered representations of shape plausibly obeys such an ecological constraint. Viewer-centered representations are computed from cues to distance and slant, such as binocular disparity, texture, and specular reflections (Fleming et al. 2004). Conversely, researchers have proposed that object-centered representations are computed independently of distance/slant cues, perhaps by applying global Gestalt constraints like symmetry and compactness directly to the 2D retinal image (Li 2009; Li & Pizlo 2011). If ecological generalizations ensure that the cues used by these two computations regularly agree concerning the most likely distal shape, then representations produced in these two ways will converge, respecting the Shape Exclusion Law. However, in environments where one type of cue is either misleading or impoverished, the resulting representations may diverge (see section 4.1).

I turn to computational constraints. By *computational constraints*, I mean constraints on patterns of information flow within a system that persist through changes in the specific information it happens to process on any given occasion. For example, it might be a computational constraint on human minds that belief cannot affect perception (Pylyshyn 1999), or that visual working memory capacity is limited to a fixed quantity of objects or resources (Bays & Husain 2008). The visual system's computational constraints are adapted to its task of recovering conditions in the environment. It possesses stable dispositions to form and update certain representations on the basis of other representations or information-carrying states. Thus, the visual system might be stably disposed to derive representations of surface convexity from registrations of luminance patterns in the retinal image, and to update these representations based on stereoscopic depth information.

If two properties are incompatible in the world, it should not be surprising to find this fact mirrored in the visual system's computational constraints. It would plausibly have been advantageous during evolution for the architecture of the visual system to ensure that we don't *see* things as both red and green, since nothing in our environment could *be* both red and green, and there is no obvious benefit to committing this type of perceptual error.

What specific kinds of computational constraints might promote conformity to the laws of appearance? One possibility is that the visual system is stably disposed to derive representations of one property from representations of another, and these computational priority relations prevent certain conflicts from arising. For example, Davies (2021, 591) suggests that the visual system might be disposed to compute representations of monadic surface color on the basis of color relations. Such computational priority might preclude certain conflicts between experienced monadic colors and color relations (for instance, experiencing *A* as red, *B* as green, but *B* as redder than *A*). Computational priority relations can also generate dependencies between representations of different dimensions, when one dimension is computed from the other or both are computed from a common basis. If the visual system is stably disposed to compute representations of an object's real-world size from its retinal size and represented distance (Gogel 1969), then it may be prevented, under normal circumstances, from representing an object's size without representing its distance. Furthermore, if both object-centered and viewer-centered shape representations are derived from integrated contour maps in early vision (Kellman & Fuchser 2023), this shared input to both computations might prevent them from exhibiting certain topological conflicts—for instance, concerning whether a shape's contour is closed or open.

The visual system is also stably disposed to perform certain *conflict-resolution* operations: When distinct representations of a given dimension, such as size or slant, are derived from separate cues, the distinct single-cue estimates are updated in a manner sensitive to the cues' reliability in order to promote coherence between them (Ernst & Banks 2002; Rescorla 2020). Accordingly, if experience reflects the outcomes of conflict-resolution operations (rather than their inputs), then inconsistencies between visual representations of a given dimension may be resolved before becoming phenomenologically manifest.

One might object that there is no principled distinction between format constraints and computational constraints. Perhaps stable dispositions toward conflict-resolution are simply the *grounds* of the incompatibility constraint, since such conflict-resolution operations ensure that syntactic properties encoding incompatible values of a dimension do not co-occur. I grant that a representation's format properties may be *among* its computational properties (or may supervene on its overall computational role), so format constraints can be reasonably regarded as a *subset* of computational constraints. Nevertheless, the class of computational constraints is broader, since computational constraints are often "softer" than format constraints. The distinction between the hybrid approach from the format-based approach is that the former recognizes the significance of such *format-extrinsic* computational constraints in explaining the laws of appearance. Let me elaborate.

If a constraint is built into the format of perceptual representation, then violations of the constraint should require the production of syntactically ill-formed visual representations, akin to ungrammatical sentences of English. It is an open question whether the visual system ever produces ill-formed representations.[8] But if so, then such representations should minimally result from some processing malfunction, and should exhibit other abnormalities. For instance, they should be incapable of participating in downstream perceptual computations in the normal way. Thus, one might expect a syntactically ill-formed representation of size to be incapable of participating in processes of recognition and categorization that take perceived size as input.

---

[8] Lande (2021) regards syntactically ill-formed representations as "psychologically impossible" (651), though I need not commit to this here. Certain format constraints may be implemented by physiological limitations of the perceptual system, such that violating the constraint would require the system to enter a state that it is physiologically unable to enter. For instance, perhaps violations of the incompatibility constraint within particular visual maps would require incompatible patterns of activity in particular pools of neurons.

I offer two examples to illustrate how computational constraints may be softer than format constraints.

First, computational priority relations can be disrupted by bringing about a representation in atypical ways without necessarily leading the resulting representation to be ill-formed. Suppose again that the visual system is disposed to derive representations of size from representations of distance, normally preventing it from representing size without distance. Still, one might *induce* the visual system to represent size without distance by directly stimulating relevant sensory areas, and the resulting size representations need not be ill-formed. They may, for instance, function in downstream computations like recognition or categorization in typical ways.

Second, computational constraints can strongly discourage the formation of a representation not because it would be ill-formed, but because it represents a scenario the system deems highly improbable. The development of cue-integration in children offers a plausible example of this situation.

Children gradually learn to integrate visual cues to slant, reaching adult-like performance only around age 12 (Dekker et al. 2015). Importantly, in adults, the integration of textural and stereoscopic slant cues is *complete*: When two stimuli differ in both stereoscopic and textural cues, but agree in the precision-weighted average of these single-cue estimates, adults are unable to distinguish them (Hillis et al. 2002; Nardini et al. 2010), suggesting that the single-cue estimates are replaced by a single, cue-independent slant representation. Before age 12, children *outperform* adults on this type of discrimination task, suggesting that they retain two representations based on texture and stereoscopic cues, respectively (Nardini et al. 2010; Dekker et al. 2015).

However, cue integration does not emerge all at once. Around 10 years of age, texture-based and stereoscopically-based slant representations *influence* one another, but adult performance has not yet been reached (Dekker et al. 2015). At this age, the child's visual system is biased such that texturally-based and stereoscopically-based slant representations reliably agree under normal circumstances, but exceptions are permitted in unusual cases. In Bayesian terms, the child's visual system internalizes a "coupling prior" expressing high but imperfect confidence that texturally- and stereoscopically-defined slant should be consistent (Ernst 2007; Rescorla 2020). The system must learn that these variables are perfectly correlated. Because the coupling prior is imperfect, it can be swamped when textural and stereoscopic cues sharply diverge. Consequently, conflicting representations are retained.

The strong but imperfect coupling prior imposes a computational constraint on the child's visual system urging coherence among slant representations. However, this constraint is not credibly regarded as a *syntactic rule*, since representations violating it needn't be ill-formed. We should posit syntactically ill-formed visual representations only given some identifiable processing failure or malfunction. But no malfunction occurs when a child's visual system retains conflicting slant estimates given abnormally large divergence between slant cues. Instead, this reflects *optimal functioning* for a system with high but imperfect coupling priors. The system represents a scenario that it deems *highly unlikely but not impossible*: i.e., a situation where texture-defined and stereoscopically-defined slant are inconsistent. Moreover, the conflicting slant representations may participate in downstream computations such as action guidance in the normal way. I see no good reason to classify such representations as ill-formed. Thus, the child's coupling prior imposes a format-extrinsic computational constraint on visual processing of slant.

Summing up, not all computational constraints are format constraints. Computational constraints are often softer than format constraints, since they can be violated without producing syntactically ill-formed representations. I suggest that such format-extrinsic computational constraints often promote coherence between distinct visual representations of a given property or dimension. Of course, given their softness, we should expect them to permit more frequent exceptions than format constraints. I return to this point in section 4.

3.2. Other Laws of Appearance

The hybrid approach promises to generalize to other putative laws of appearance. Recall Berkeley's Law and the Perspectival Law:

> *Berkeley's law.* (1) An individual cannot experientially represent that something has a color without also experientially representing that it takes up space in some way. (2) Conversely, an individual cannot experientially represent that something takes up space in some way (e.g., being circular) without also experientially representing a qualitative difference (e.g., a *white* circle on a *black* background).
>
> *Perspectival law.* An individual cannot experientially represent merely *that there is a cube somewhere in reality*, without any "perspectival content" about its location and apparent shape from "here." (Pautz 2021, 131)

Both laws embody constraints in which experiences of certain properties *require* or *depend upon* experiences of other properties. The first version of Berkeley's law states that the experience of color requires the experience of spatial location, while the second states that the experience of shape requires the experience of some qualitative difference between the shape's interior and exterior. The Perspectival Law involves two dependence relations: first, the experience of shape requires the experience of egocentric location, and second, experience of *intrinsic* shape requires experience of *perspectival* shape (e.g., whenever you experience a slanted penny as intrinsically circular, you must also experience it as 'elliptical-from-here').[9]

I suspect that Necessitism is less intuitively compelling for these laws than for the Exclusion Laws, and Pautz seems to agree. Indeed, as Pautz observes, there is evidence that Berkeley's Law is contingent, since the law is plausibly violated in individuals with cerebral achromatopsia (Pautz 2021, 133). But assuming that these laws obtain at least under normal circumstances, perception science

[9] For discussions of perspectival shape properties, see Green & Schellenberg (2018) and Lande (2018).

offers ample resources for explaining them. Dependencies between perceptual representations of distinct properties may derive from format, computational, or ecological constraints, or combinations thereof. It is an empirical question how these constraints interact in any given case. But just to illustrate, consider the first dependency laid out in the Perspectival Law: experiential representation of shape requires experiential representation of egocentric location. Fortunately, the picture of shape experience developed here demonstrates how all three constraints might fruitfully combine in explaining this law.

I've suggested that shape is simultaneously experienced in a viewer-centered manner and an object-centered manner. Plausibly, it is a format constraint on *viewer-centered* shape representation that well-formed representations of this type have representations of direction, and perhaps distance, as constituents (Lande 2018). Furthermore, specifying an object's distance and direction suffices to specify its egocentric location. Thus, format constraints may explain why an object's shape cannot be experienced in a *viewer-centered manner* without experiencing its egocentric location.

However, object-centered representations obey no similar format constraint, since they do not specify shape in terms of egocentric distance and direction. Indeed, a signature virtue of object-centered representations is that they remain invariant through *changes* in egocentric location. Nonetheless, format-extrinsic computational or ecological constraints may ensure that object-centered representations virtually never occur without viewer-centered representations. Perhaps there is a computational constraint whereby the visual system is stably disposed to derive object-centered representations from viewer-centered representations (Marr 1982; although see Bennett 2012). Alternatively, it might be an ecological constraint that any object that supplies proximal cues sufficient to prompt construction of an object-centered shape representation also supplies cues sufficient to prompt construction of a viewer-centered representation. Further, because the visual system is stably disposed to form representations of both types whenever the relevant cues are available, it embodies the regularity that whenever shape is represented in *either* a viewer-centered *or* object-centered manner, location is represented as well.

This explanation of the Perspectival Law inspires testable hypotheses about when the law may be violated. If the processes that produce viewer-centered shape representations can be selectively impaired, then a person might end up forming object-centered representations without viewer-centered representations, perhaps even experiencing an object's shape without experiencing its location. Just this phenomenon has been claimed to occur in Bálint's syndrome, where patients reliably identify the shapes of objects but struggle to judge both the absolute and relative locations of objects (for example, whether they are in the right or left hemifield; Friedman-Hill et al. 1995; Schwenkler 2012; Roberston 2014; see also Vannuscorps et al. 2022). Perhaps these cases occasionally involve visual experiences of shape without location. At least, this matches the way such patients *describe* their experiences (Robertson 2004, 158-159; Schwenkler 2012, 316).[10]

3.3. Do Format Constraints Ever *Suffice?*

I've argued that format constraints do not suffice to explain the Exclusion Laws, though they may enter into satisfactory explanations of these laws. However, I grant that format might suffice to explain *certain* experiential regularities. I offer two examples.

First, if there are experiential regularities entirely subserved by a single type of representational structure, then format may sometimes suffice to explain them. In particular, format constraints might ensure that a property or dimension cannot be represented *in a particular way* without also representing certain related properties or dimensions. Perhaps we cannot experience shape in a viewer-centered manner without also experiencing distance and direction because

[10] See, however, French (2018) for an alternative view on which Bálint's patients experience objects' locations, but are merely limited to experiencing the location of one object at a time.

representations of the latter properties are constituents of any well-formed viewer-centered shape representation. Furthermore, if there are sensory modalities (or subsystems of sensory modalities, like vision-for-action) that represent a given dimension via a single type of representational structure, then it may be possible to explain variants of the Exclusion Law for that dimension and sensory modality (or subsystem) solely by appeal to format. I doubt that this is the case for the Color or Shape Exclusion Laws, given the heterogeneity of our visual representations of color and shape, but it is an empirical question whether the situation occurs for other dimensions or modalities. If it does, then representationalists might employ format-based explanations to explain such variants of the Exclusion Law.

Second, even when a property is visually represented by multiple representational structures, all of them might happen to obey common format constraints. In such cases, format constraints within each structure might also apply more generally. Perhaps all extant systems of perceptual shape representation require well-formed shape representations to be composed of representations of at least some spatial parts of a shape, such as edges, curvature segments, or surface patches (Elder 2018, 436; Ashby 2022). If so, format constraints alone might explain why one cannot experientially represent a medium-sized object's shape without also experientially representing spatial properties of certain of its parts.[11]

However, the fact that format *sometimes* suffices to explain experiential regularities is not problematic for the hybrid approach, which aims to capture the *general form* that explanations of the laws of appearance should adopt. The relative significance of the three factors is expected to vary depending on the specific experiential regularity under consideration, so we should view these

[11] The present view is also compatible with Ashby's (2022) proposal that format constraints on shape representation might preclude "shape inversion" scenarios in which experiences phenomenally like our experiences of spheres serve to represent cubes, and vice versa (Ashby 2022: 399)). Ashby argues that such scenarios are ruled out by the requirement of a systematic one-to-one relationship between the constituents of a shape representation and the properties they represent (Ashby 2022: 400-401). Perhaps all viable systems of shape representation obey such a format constraint.

examples as limiting cases where format reaches maximal explanatory significance while the other factors reach minimal significance. We might also expect cases in which an experiential regularity is fully explained by either ecological or computational constraints without appeal to format.

Summing up: I have argued that format constraints do not suffice to explain the laws of appearance—in particular, they do not suffice to explain the Exclusion Laws. Thus, the modified format explanation fails. In its place, I have recommended a hybrid approach on which format-extrinsic computational and ecological constraints also contribute to explaining the laws of appearance. The latter constraints help explain how coherence is maintained not just *within* representational structures, but *between* them, remedying a critical flaw in the modified format explanation.

4. The Challenge from Necessitism

However, the hybrid approach faces a challenge. As Pautz observes, our pretheoretic intuitions favor Necessitism, according to which the Exclusion Laws are metaphysically necessary. But if the hybrid approach is correct, then it is likely that these laws are contingent. Thus, unless the contingency of the Exclusion Laws can be defended, we should reject the hybrid approach.

I now argue that the Exclusion Laws *are* plausibly contingent, since there is evidence for real-world violations of the Shape Exclusion Law. Thus, we should have serious doubts about Necessitism. So the fact that the hybrid approach conflicts with Necessitism is not a good reason to reject it.

Recall that empirical explanations of the laws of appearance are potentially compatible with either Necessitism or Contingentism, depending on whether the facts figuring in the explanation constitute essential or accidental features of perception. The hybrid approach identifies three types of facts that conspire to explain the laws of appearance: format constraints, ecological constraints, and computational constraints. Could any of these facts constitute essential features of perception?

Some have claimed that certain format features are indeed essential. Block (2023, 24-25, 50) holds that perception is *necessarily* iconic. If a system transformed patterns of proximal stimulation directly into discursive representations, then it would not be perceiving. While I am skeptical of this claim (Green 2023b, 471, 487), I will not attempt to adjudicate the issue here because the challenge to the hybrid approach arises even Block is correct.

Even if certain format constraints are essential to perception, many ecological and computational constraints plainly are not. Indeed, we might expect such constraints to fail in real-world cases. The environmental generalizations that place ecological constraints on perception—for instance, correlations between perceptual cues to the same property—should have exceptions. And computational constraints that promote conflict-resolution among visual representations in normal conditions might fail to produce coherence for unusual stimuli. Analogously, while the architecture of opponent-processing mechanisms might ensure that things do not appear reddish-green under normal circumstances, unusual stimuli can appear this way (Billock & Tsou 2010).

Thus, if the laws of appearance rest partly on ecological or computational constraints, then we should expect these laws to be contingent. However, this verdict conflicts with the pretheoretical intuition that they are necessary.

We have already encountered some evidence against Necessitism. I suggested earlier that the Perspectival Law may be violated in individuals with Bálint's syndrome, while Pautz grants that Berkeley's Law may be violated in individuals with cerebral achromatopsia. Accordingly, Pautz rests his case for Necessitism almost entirely on the Exclusion Laws. Regarding the Shape Exclusion Law, he contends: "Unlike in the case of Berkeley's law, we have no reason to believe that there are actual cases where it is violated" (2021, 133-134).

Contra Pautz, I argue that there are real-world cases in which the Shape Exclusion Law is violated. Thus, although the hybrid approach suggests that the law is contingent, that implication is defensible. Moreover, once we understand how its violations arise, we can understand why we remain tempted to think that the law is necessary.

I discuss two strands of evidence for violations of the Shape Exclusion Law. The first involves inconsistencies between object-centered and viewer-centered representations of shape. The second involves inconsistencies among perceptual representations of slant, relative depth, and curvature within a single surface.

4.1. Violations of the Shape Exclusion Law

4.1.1. Object-Centered vs. Viewer-Centered Representation. I have argued that visual experience is subserved by both viewer-centered and object-centered shape representations, and that these representations are formed based on distinct cues and computational principles. Reliable coherence between the representations is facilitated by correlations among the cues used to construct them. However, these correlations can be broken in impoverished stimuli that contain poor cues to distance and surface orientation, but strong cues to object-centered shape representation.

In a test of consistency between viewer-centered and object-centered representation, Li (2009) presented line drawings of polyhedra that elicited 3D percepts (fig. 3a). Participants had two tasks. First, they adjusted the shape of an elliptical probe (fig. 3b) until it appeared as a circle lying on each visible face of the polyhedron. Second, they adjusted a parallelogram (fig. 3c) until it appeared the same intrinsic shape as each visible face of the polyhedron. The first task plausibly taps viewer-centered representation, since it assesses the perceived orientation of the parallelogram's visible surfaces. The second plausibly taps object-centered representation, since it involves generalizing between the rotated polyhedron and head-on parallelogram. Li constructed two shapes based on subjects' responses to determine whether the viewer-centered and object-centered representations used for the tasks were consistent. They were not: Responses in the parallelogram-adjustment task showed a much stronger tendency toward rectangular faces than did elliptical-probe responses (fig. 4). Such discrepancies likely arise because object-centered shape representation is more influenced by Gestalt principles (e.g., a bias toward global symmetry) relative to viewer-centered representation.

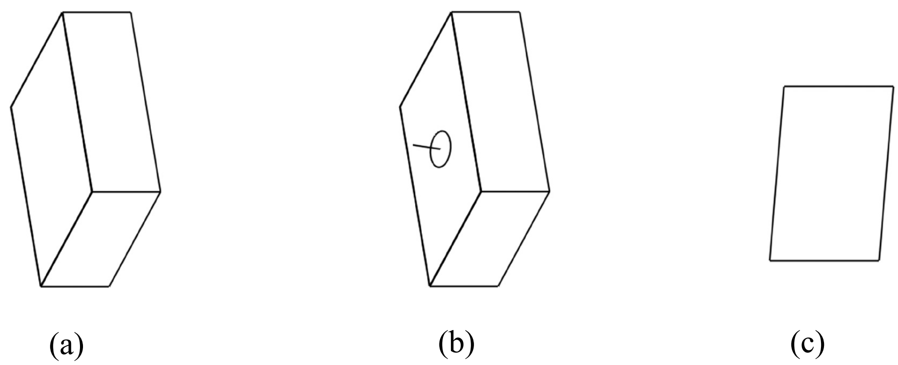

Figure 3. (a) Sample polyhedron from Li (2009). (b) The elliptical-probe task. (c) The parallelogram-adjustment task.

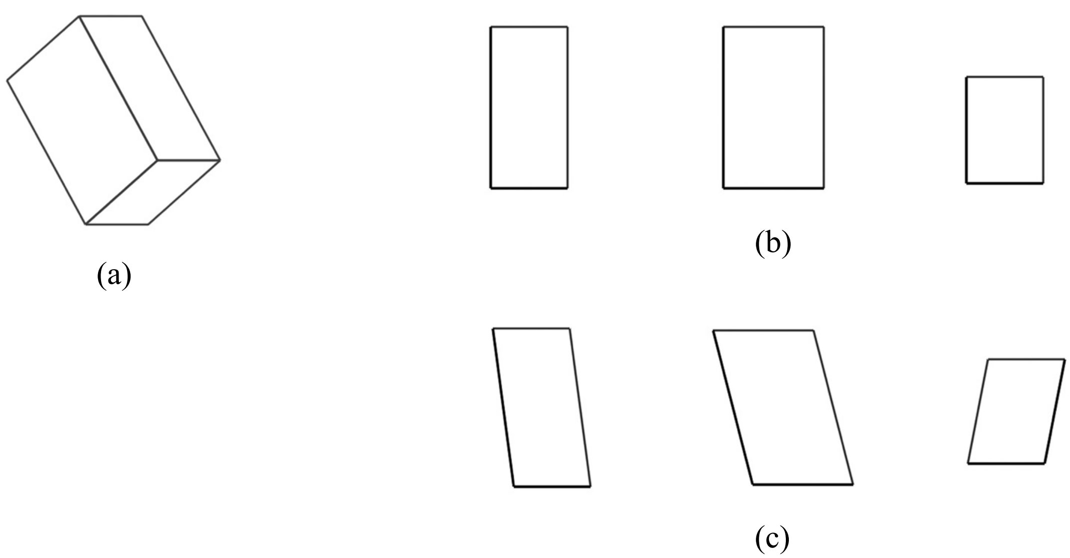

Figure 4. (a) A polyhedron from Li (2009). (b) Faces of the polyhedron constructed from subjects' responses in the parallelogram-adjustment task. (c) Faces constructed from subjects' responses in the elliptical-probe task.

While no single study is decisive, these results strongly suggest that viewer-centered and object-centered shape representations can diverge for a single object at a time (see also Li & Pizlo 2006). If both representations contribute to the overall visual experience of shape, as argued earlier, then we have a violation of the Shape Exclusion Law. Note that these violations are fairly subtle, concerning the precise angles of an object's faces—not on the magnitude of seeing an object as both square and circular. But they are Exclusion-Law violations all the same.

One might suggest that participants never enjoyed conflicting experiences of viewer-centered and object-centered shape contents *simultaneously*, but rather underwent the two experiences in succession, making judgments on the basis of one or the other experience (compare Bayne 2010, 54). If so, however, one would expect to find phenomenological signatures of these shifts when viewing the shapes in figures 3 and 4. Indeed, one might expect the sort of Gestalt shifts that occur elsewhere when different reference frames can be imposed on a figure and we alternate between them—e.g., when a figure can be seen as either a diamond or a tilted square. However, I ascertain no such shifts when viewing the stimuli.

Alternatively, one might suggest that responses in the parallelogram-adjustment task did not tap into online experience. Perhaps people mentally rotated the slanted polyhedron until the relevant face was viewed head-on, then reported the shape of that face; and perhaps the bias toward rectangularity was introduced during mental rotation. While this is an empirical possibility, I find it unlikely. When I view the shape in figure 4a, I find it hard to deny that the apparent rectangularity of its faces is immediately present in visual experience. The shape simply *looks like* a box composed of rectangular faces. Nonetheless, these issues illustrate why I regard the evidence as highly suggestive, but not decisive. So it is helpful to have further evidence for violations of the Shape Exclusion Law.

4.1.2. Representations of Curvature, Slant, and Depth. I turn to a second case. There is evidence that perceptual representations of various aspects of surface geometry—namely, curvature, slant, and relative depth—can be mutually inconsistent: The apparent curvature of a surface region may be inconsistent with the apparent slant and depth of that same region, violating the Shape Exclusion Law.

Di Luca et al. (2010) showed participants pairs of paraboloid shapes like those in figure 5 (top), which typically appear as curved surfaces protruding toward the subject. On each trial, one shape was held constant while the other was presented at varying elongations. Participants had to choose the elongation of the adjustable shape that best matched the constant shape in either (a) the apparent curvature at its tip, (b) the apparent slant of a surface patch within the shape (i.e., how far its tangent plane was rotated out of the frontoparallel plane), or (c) the apparent depth difference between the tip of the paraboloid and another point on the surface.

Note that when a paraboloid is elongated, this increases the magnitude of (a)-(c) simultaneously (see figure 5, bottom). Accordingly, the paraboloid that appears more slanted should *also* appear more curved, and so on. Surprisingly, however, di Luca et al. found that comparative judgments about the three properties were mutually inconsistent. Certain of their judgments required one paraboloid to be longer than the other, while others required the reverse. Thus, regarding the shapes in figure 5 (top), di Luca et al. write:

> Typically, the curvature at the center of the display is perceived to be larger for the shading stimulus. For a local curvature-match, therefore, we should reduce the elongation of the shading stimulus. (…) [H]owever, observers perceive a larger relative depth between the tip of the paraboloid and the bounding-contour for the texture than for the shading rendering. For a depth-match, therefore, we should reduce the elongation of the texture stimulus. Finally, near the boundary of the display, local slant appears to be shallower for the shading stimulus. For a local slant-match, therefore, we should increase the elongation of the shading stimulus. This example shows that the perceptual judgments of depth, slant, and curvature are mutually inconsistent. (di Luca et al. 2010: 1527)

Assuming that perceptual experience represents each of these aspects of surface geometry, we have another apparent violation of the Shape Exclusion Law.[12] Moreover, such violations were not limited to cases where the shapes were defined by a single cue, but generalized to shapes defined by multiple cues concurrently (such as texture, shading, and motion).

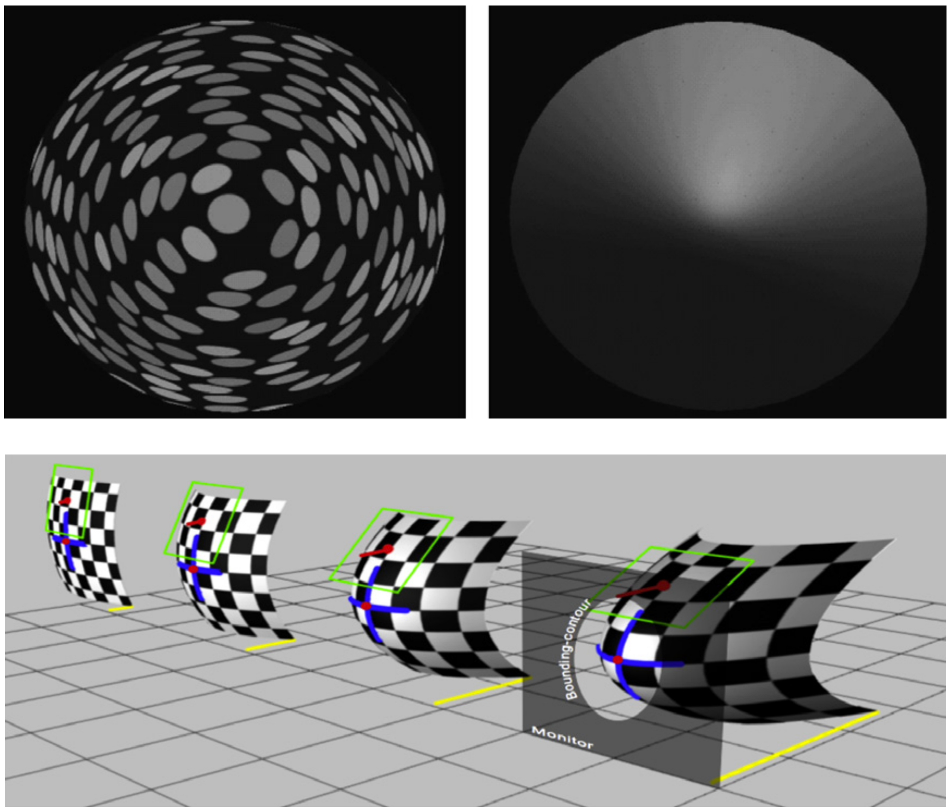

Figure 5. *Top:* Sample stimuli from di Luca et al. (2010). *Bottom:* Illustration of how paraboloid elongation simultaneously increases curvature, slant, and relative depth. In the rightmost shape, the tip of the paraboloid is most curved and the surface patch indicated by the green square is most slanted. The depth difference between the tip and a second surface point is also largest (indicated by the red segment).

We can explain these puzzling results on the assumption that curvature, slant, and depth are perceptually recovered through distinct computations capitalizing on distinct cues in the proximal stimulus, and that these computations generate distinct representations that each suffice to specify surface shape (e.g., a "depth map," "slant map," and "curvature map" (see Di Luca et al. 2010,

[12] For further evidence challenging the internal consistency of our experiences of surface geometry, see Domini et al. (1998), discussed in Domini and Caudek (2003).

1521). Accordingly, one paraboloid might be represented as more elongated than another within their respective depth maps, but as less elongated within the curvature maps, explaining inconsistencies in participants' comparative judgments. Presumably, *radical* conflicts among these representations rarely if ever occur, but subtle conflicts sometimes arise.

I have discussed two examples where the Shape Exclusion Law seems to be violated. The reason for both violations is fundamentally the same: Shape is represented by the visual system in multiple ways concurrently, based on distinct cues and computations. Stable ecological and computational constraints ensure that these representations regularly agree. However, these constraints are not foolproof, so occasionally we get violations of the Shape Exclusion Law. Thus, we should have serious doubts about Necessitism.

4.2. Diagnosing Necessitism's Appeal

Why, then, are we pretheoretically inclined to *believe* that the Shape Exclusion Law is inviolable? Here is a tentative conjecture: (i) While the Shape Exclusion Law has counterexamples, these are not easily detectable by introspection. Thus, we *take* our introspective evidence to indicate that every experience we have undergone obeys the law. (ii) Because we also assume that our shape experiences form a comprehensive, unbiased sample of the phenomenal kind to which they belong, we infer that all *possible* shape experiences obey the Shape Exclusion Law.

By a "phenomenal kind," I mean, roughly, a family of phenomenal properties unified by relations of similarity, difference, and indistinguishability (see Rosenthal 2010; Lee 2021). We can assess whether the phenomenal character of seeing a square is more like that of seeing a rectangle or that of seeing a circle. Furthermore, it is possible to move from experiences of squares to experiences of either rectangles or circles via a series of experiences where adjacent experiences in the series are either fully indistinguishable or at least highly confusable. These data suggest that the three phenomenal properties belong to the same phenomenal kind. Conversely, we cannot assess whether the phenomenal character of seeing square is more like that of seeing redness or that of seeing blueness, and we cannot move between these experiences through a series of near-indistinguishable steps, suggesting that they do not all belong to a single phenomenal kind. If there are experiences of "alien" colors incommensurate with the color properties we perceive, they would belong to a distinct phenomenal kind from our color experiences. However, because Exclusion-Law violations involve appearances of the *same* shape or color properties we are already familiar with, arguably they should belong to the same phenomenal kinds as familiar shape or color experiences.

Now, regarding (i): Shape is a complex property that can be represented in various ways using different primitives, combinatorial rules, and reference frames (Elder 2018; Green 2023; Lande 2024). Suppose, then, that perception generates two shape representations that differ in these ways, and experience also represents shape in these two ways concurrently. Small inconsistencies between the properties so represented might be computationally difficult to recover, and obscured from introspective detection. If consistency is always maintained *within* each representation (perhaps through more robust format constraints), then we might be led to believe that all our shape experiences obey the Exclusion Law.

Regarding (ii): If we *take* our introspective evidence to indicate that every shape experience we have undergone obeys the Exclusion Law, then we might use familiar heuristics to infer that all *possible* shape experiences must obey the Exclusion Law.

Consider an analogy from Nichols et al. (2016, 533). Suppose Mary has two dice with denominations 6 and 10. She rolls one of them 100 times, and each result falls between 1 and 6. If you were asked which die Mary rolled, you would choose the 6-sided die. After all, if she had rolled the 10-sided die, and the dice are unbiased, you would expect *some* of the rolls to come up greater than 6. Likewise, suppose we assume that our shape experiences form an unbiased sample from the phenomenal kind *visual shape experience*, and we believe (mistakenly) that none of our experiences violate the Shape Exclusion Law. Then, if visual shape experiences *could* violate the Exclusion Law, the introspective datum would seem like a suspicious coincidence. For if shape experiences *can* violate the Exclusion Law, why do we evidently never have such experiences? Thus, we might favor the hypothesis that *no* shape experiences can violate the Exclusion Law.

Pautz observes that the reason we are attracted to Necessitism about the Exclusion Laws cannot merely be that we are unable to imagine experiences that violate them. Arguably we cannot imagine experiences of bat echolocation, but we readily grant their possibility. However, if I am right, our attraction to Necessitism is *not* just based on our inability to imagine counterexamples to the Exclusion Laws, but also on the assumption that our experiences are representative, unbiased samples of the phenomenal kinds they instantiate. Because we do *not* think that echolocative experiences instantiate the same phenomenal kind as any of our experiences, we have no parallel basis to conclude that echolocative experiences are impossible.

One might object that this account overgeneralizes because it predicts that certain hypothetical laws of appearance *should* strike us as necessary when in fact they do not. Consider a hypothetical *Cross-Modal Shape Exclusion Law*: One cannot visually experience an object as having one shape while haptically experiencing it as having a different shape. While conflicting visual and haptic experiences of shape strike us as odd or unusual, and our perceptual systems seek to resolve them (Helbig & Ernst 2007), few would be tempted to say that such conflicts are *impossible*.

However, one might think that my explanation of the apparent necessity of the standard Shape Exclusion Law should generalize to the Cross-Modal Exclusion Law. For, assuming that shape is coded differently in separate modalities, conflicts between visual and haptic shape representations should be difficult to detect introspectively, just like viewer-centered/object-centered conflicts. Now consider the range of *multisensory shape experiences* we have undergone: that is, cases of both visually and haptically experiencing an object's shape. If we are not aware of any cases in which multisensory shape experience has violated the Cross-Modal Exclusion Law, *and* we assume that our past multisensory shape experiences constitute a representative sample of the kind *multisensory shape experience*, then we should be prepared to believe that the Cross-Modal Shape Exclusion Law is inviolable. Since we do not believe this, we must have some stronger basis for believing that the unimodal Shape Exclusion Law is inviolable.

I reply that the cases are not truly analogous. While violations of the Unimodal Shape Exclusion Law are both atypical and relatively small, most perceivers are familiar with large, salient violations of its cross-modal counterpart (see also Bayne 2010, 56-57). This stems from the fact that haptic systems recover shape through bodily contact, not optical projection, enabling shape features that are unrecoverable from the retinal image to be perceived by touch. Suppose you see a coffee mug that visually appears cylindrical. However, you grasp it and discover that its backside is entirely flat. Here, apparent shape differs radically across modalities. You visually experience the object as cylindrical, while haptically experiencing it as a half-cylinder. I suggest that while subtle conflicts within vision tend to elude introspective detection, significant cross-modal conflicts like this are both familiar and easily noticed. Accordingly, unlike the unimodal case, we *are* aware of (or implicitly sensitive to) examples where multisensory shape experience violates the Cross-Modal Shape Exclusion Law, so we are not tempted to think that it is necessary.

I have argued that while we are strongly tempted to believe the Shape Exclusion Law is necessary, it is actually contingent. If so, we should be suspicious of Necessitism regarding other coherence principles about perception. The Shape Exclusion Law seems just as plausible as any other coherence principle, so if *it* is contingent, others could easily turn out to be contingent as well—and for similar reasons. Specifically, violations of coherence principles are most likely when a single property is perceptually represented in multiple ways based on schemes that specify the property differently. However, such differences should *also* make coherence violations particularly difficult for us to detect introspectively. Furthermore, if we assume that our perceptual experiences provide representative samples of the phenomenal kinds they instantiate, we might reasonably believe that if there *were* experiences violating the relevant coherence principles, we should have undergone them. Because we think (mistakenly) that we have not undergone them, we conclude (mistakenly) that the principles are inviolable.

Thus, I propose that we can employ evidence from cognitive science not only to explain why the laws of appearance generally hold, but also to pinpoint the conditions under which they are violated, and why, despite these violations, we might continue to think that they are necessary.

5. Conclusion

It is widely thought that there are limits on the ways that things can appear to us in perceptual experience, or "laws of appearance." Assuming that such laws obtain, and that they are not brute, inexplicable facts, how are they to be explained? One option is that they are explained by a priori knowable facts about perceptual experience. Another option is that they are explained by empirically discoverable features of our perceptual systems. This question carries special significance given that leading contemporary theories of experience, such as representationalism, do not support a priori explanations of the laws. Such theories are in need of a viable empirical explanation.

This paper has outlined the form that I believe an empirical explanation of the laws of appearance should adopt. While some authors propose that the laws result purely from the format of perceptual representation, I advocate a hybrid approach, which appeals not merely to format, but to ecological and format-extrinsic computational constraints on our perceptual systems. While the hybrid approach suggests that the laws are contingent, I have defended this implication by adducing evidence for counterexamples to the Shape Exclusion Law. In fact, this evidence also poses a challenge to philosophers who favor a priori explanations on which the laws are necessary: If perception science indicates that the laws are *not* necessary, then we should not accept an explanation which entails that they are.

More generally, I suggest that our pretheoretical convictions about which kinds of experiences are possible should not be taken at face value, but neither should they be flatly dismissed. The laws of appearance are not mere accidents, but rather lawlike generalizations explained in terms of stable properties of our environments and perceptual systems. These generalizations are discoverable through empirical investigation. The hybrid approach highlights the potential of perception science both to explain why the laws of appearance hold under ordinary circumstances, and also to elucidate the conditions where they can be violated.

References

Akins, K. A., & Hahn, M. (2014). More than mere colouring: The role of spectral information in human vision. *The British Journal for the Philosophy of Science*, *65*(1), 125-171.

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

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*(5890), 851-854.

Beck, J. (2019). Perception is analog: The argument from Weber's Law. *The Journal of Philosophy*, *116*(6), 319-349.

Bennett, D. J. (2012). Seeing shape: Shape appearances and shape constancy. *The British Journal for the Philosophy of Science*, *63*(3), 487-518.

Berkeley, G. (1713/1906). *Three Dialogues Between Hylas and Philonous*. Chicago, IL: Open Court.

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.

Billock, V. A., & Tsou, B. H. (2010). Seeing forbidden colors. *Scientific American*, *302*(2), 72-77.

Block, N. (2023). *The Border Between Seeing and Thinking*. Oxford: Oxford University Press.

Briscoe, R. E. (2008). Vision, action, and make-perceive. *Mind & Language*, *23*(4), 457-497.

Broad, C. D. (1925). *The Mind and its Place in Nature*. London: Kegan Paul.

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

Camp, E. (2018). Why maps are not propositional. In A. Grzankowski & M. Montague (eds.), *Non-Propositional Intentionality*, 19-45. 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.

Davies, W. (2021). Colour relations in form. *Philosophy and Phenomenological Research*, *102*(3), 574-594.

Dekker, T. M., Ban, H., van der Velde, B., Sereno, M. I., Welchman, A. E., & Nardini, M. (2015). Late development of cue integration is linked to sensory fusion in cortex. *Current Biology*, *25*(21), 2856-2861.
Di Luca, M., Domini, F., & Caudek, C. (2010). Inconsistency of perceived 3D shape. *Vision Research*, *50*(16), 1519-1531.
Domini, F., & Caudek, C. (2003). 3-D structure perceived from dynamic information: A new theory. *Trends in Cognitive Sciences*, *7*(10), 444-449.
Domini, F., Caudek, C., & Richman, S. (1998). Distortions of depth-order relations and parallelism in structure from motion. *Perception & Psychophysics*, *60*, 1164-1174.
Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
Elder, J. H. (2018). Shape from contour: Computation and representation. *Annual Review of Vision Science*, *4*, 423-450.
Erdogan, G., & Jacobs, R. A. (2017). Visual shape perception as Bayesian inference of 3D object-centered shape representations. *Psychological Review*, *124*(6), 740-761.
Ernst, M. O. (2007). Learning to integrate arbitrary signals from vision and touch. *Journal of Vision*, *7*(5):7, 1-14.
Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429-433.
Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences*, *103*(47), 18014-18019.
French, C. (2018). Bálint's syndrome, object seeing, and spatial perception. *Mind & Language*, *33*(3), 221-241.
Green, E. J. (2017). A layered view of shape perception. *The British Journal for the Philosophy of Science,* *68*(2), 355-387.
Green, E. J. (2019). On the perception of structure. *Noûs*, *53*(3), 564-592.
Green, E. J. (2022). 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.
Green, E. J., & Schellenberg, S. (2018). Spatial perception: The perspectival aspect of perception. *Philosophy Compass*, *13*(2), e12472.
Gregory, E., & McCloskey, M. (2010). Mirror-image confusions: Implications for representation and processing of object orientation. *Cognition*, *116*(1), 110-129.
Haugeland, J. (1991). Representational genera. In W. Ramsey, S. Stich, & D. E. Rumelhart (eds.), *Philosophy and Connectionist Theory*, pp. 61-89. Hillsdale, NJ: Erlbaum.
Helbig, H. B., & Ernst, M. O. (2007). Optimal integration of shape information from vision and touch. *Experimental Brain Research*, *179*, 595-606.
Hillis, J. M., Ernst, M. O., Banks, M. S., & Landy, M. S. (2002). Combining sensory information: Mandatory fusion within, but not between, senses. *Science*, *298*(5598), 1627-1630.
Humphreys, G. W. (2016). Feature confirmation in object perception: Feature integration theory 26 years on from the Treisman Bartlett lecture. *The Quarterly Journal of Experimental Psychology*, *69*(10), 1910-1940.
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.
Kalderon, M. E. (2007). Color pluralism. *The Philosophical Review*, *116*(4), 563-601.
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.

Kentridge, R. W., Heywood, C. A., & Cowey, A. (2004). Chromatic edges, surfaces and constancies in cerebral achromatopsia. *Neuropsychologia*, *42*(6), 821-830.
Kulvicki, J. (2015). Maps, pictures, and predication. *Ergo*, *2*(7), 149-173.
Lande, K. J. (2018). The perspectival character of perception. *The Journal of Philosophy*, *115*(4), 187-214.
Lande, K. J. (2021). Mental structures. *Noûs*, *55*(3), 649-677.
Lande, K. J. (2024). Compositionality in perception: A framework. *WIREs Cognitive Science*. DOI: 10.1002/wcs.1691.
Lee, A. Y. (2021). Modeling mental qualities. *Philosophical Review*, *130*(2), 263-298.
Lee, Y.-L., Lind, M., Bingham, N., & Bingham, G. P. (2012). Object recognition using metric shape. *Vision Research*, *69*, 23-31.
Lescroart, M. D., & Biederman, I. (2013). Cortical representation of medial axis structure. *Cerebral Cortex*, *23*(3), 629-637.
Li, Y. (2009). Perception of parallelepipeds: Perkins's law. *Perception*, *38*(12), 1767-1781.
Li, Y., & Pizlo, Z. (2006). Is viewer-centered representation necessary for 3D shape perception? *Journal of Vision*, *6*(6), 268.
Li, Y., & Pizlo, Z. (2011). Depth cues versus the simplicity principle in 3D shape perception. *Topics in Cognitive Science*, *3*(4), 667-685.
Lowet, A. S., Firestone, C., & Scholl, B. J. (2018). Seeing structure: Shape skeletons modulate perceived similarity. *Attention, Perception, & Psychophysics*, *80*(5), 1278-1289.
Maley, C. J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, *155*(1), 117-131.
Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman.
Martinez, M., & Nany, B. (forthcoming). Many-to-one intentionalism. *Journal of Philosophy*.
McCloskey, M., Valtonen, J., & Cohen Sherman, J. (2006). Representing orientation: A coordinate-system hypothesis and evidence from developmental deficits. *Cognitive Neuropsychology*, *23*(5), 680-713.
Morgan, J. B. (2023). What the senses cannot 'say'. *Philosophical Quarterly*, *73*(2), 557-579.
Nardini, M., Bedford, R., & Mareschal, D. (2010). Fusion of visual cues is not mandatory in children. *Proceedings of the National Academy of Sciences*, *107*(39), 17041-17046.
Nichols, S., Kumar, S., Lopez, T., Ayars, A., & Chan, H. Y. (2016). Rational learners and moral rules. *Mind & Language*, *31*(5), 530-554.
Pautz, A. (2017). Experiences are representations: An empirical argument. In B. Nanay (ed.), *Current Debates in Philosophy of Perception*, 23-42. New York: Routledge.
Pautz, A. (2020). The puzzle of the laws of appearance. *Philosophical Issues*, *30*(1), 257-272.
Pautz, A. (2021). *Perception*. New York: Routledge.
Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: MIT Press.
Peacocke, C. (2019). *The Primacy of Metaphysics*. Oxford: Oxford University Press.
Pylyshyn, Z. W. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*(3), 341-365.
Quilty-Dunn, J. (2023). Sensory binding without sensory individuals. In A. Mroczko-Wasowicz & R. Grush (eds.), *Sensory Individuals: Unimodal and Multimodal Perspectives*. Oxford: Oxford University Press.
Quilty-Dunn, J., Porot, N., & Mandelbaum, E. (2023). The best game in town: The reemergence of the language-of-thought hypothesis across the cognitive sciences. *Behavioral and Brain Sciences*, *46*, e261.
Rescorla, M. (2020). Perceptual co-reference. *Review of Philosophy and Psychology*, *11*(3), 569-589.
Robertson, L. (2004). *Space, Objects, Minds, and Brains*. New York: Psychology Press.

Robertson, L. (2014). Balint's syndrome and the study of attention. In A. C. Nobre & S. Kastner (eds.), *The Oxford Handbook of Attention*, pp. 1062-1080. Oxford: Oxford University Press.
Rosenthal, D. (2010). How to think about mental qualities. *Philosophical Issues*, *20*, 368-393.
Sainsbury, M. (2023). Visual experience and the laws of appearance. *Erkenntnis*, *88*(7), 2933-2940.
Schwenkler, J. (2012). Does visual spatial awareness require the visual awareness of space? *Mind & Language*, *27*(3), 308-329.
Siewert, C. (2006). Is the appearance of shape protean? *Psyche*, *12*(3), 1-16.
Todd, J. T., & Petrov, A. A. (2022). The many facets of shape. *Journal of Vision*, *22*(1):1, 1-30.
Treisman, A. (1988). Features and objects: The fourteenth Bartlett memorial lecture. *The Quarterly Journal of Experimental Psychology*, *40*(2), 201-237.
Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
Tye, M. (2020). Filling in and the nature of visual experience. *The Harvard Review of Philosophy*, *27*, 59-69.
Vannuscorps, G., Galaburda, A., & Caramazza, A. (2022). Shape-centered representations of bounded regions of space mediate the perception of objects. *Cognitive Neuropsychology*, *39*(1-2), 1-50.