For PPR symposium on Ned Block's *The Border Between Seeing and Thinking*

Iconicity, 2nd-Order Isomorphism, and Perceptual Categorization

Steven Gross
Johns Hopkins University

1. Introduction

Ned Block's project is to articulate and defend a border between perception and cognition by drawing upon advances in perception science. He draws upon perception science, first, in identifying empirical signatures of perception (adaptation, rivalry, pop-out, and illusory contours), and, second, in proposing more fundamental features of perception (non-conceptuality, non-propositionality, and iconicity). The more fundamental features would explain why the signatures tend to cluster as they do and thus constitute perception as a natural kind. My comments focus on iconicity as a fundamental, constitutive feature of perception. Block's conception of iconicity requires that perceptual representations and what they represent come in degrees (or, are graded, as I'll also put it). But there are apparent cases of perception that are not obviously iconic in this sense—for example, perceptual categorizations. Various readers of Block's book have independently noted this *prima facie* tension, including Green (2023), Firestone and Phillips (2023), and Beck (2023). This commentary develops the issue and explores possible replies. The replies differ in their burdens and commitments, or what might be costs. I look forward to learning which, if any, Block might favor.

2. Block on Iconicity

Iconic representations are often said to be picture-like or image-like. But in what sense? There are various conceptions of iconicity in the literature. It's useful to put to one side a prominent conception that Block does *not* prefer. On this view (refinements aside), a representation is iconic just in case parts of the representation represent a part of the what the whole representation represents.[1] For example, if you photograph a scene and cut out a section of the photograph, generally the part you've cut out itself represents part of the scene. Not so for such non-iconic representations as sentence-tokens.

Block allows that many perceptual representations are in accord with this conception, or something like it. But he argues that not *all* are. For example, he holds that there are perceptual analogue magnitude representations that fail to have parts corresponding to parts of the magnitude they represent. He borrows an illustration from Peacocke (2019, p. 58):

> Suppose that firing rates of certain neural circuits represent time duration. Although one duration has another duration as a part, a firing rate of 50 times per second does not have a firing rate of 17 times per second as a part. (Block 2023, p. 225)[2]

[1] Versions of this Parts Principle are propounded by Kosslyn (1980), Fodor (2007), Carey (2009), Green and Quilty-Dunn (2017), and others. Refinements are needed to handle very small parts, damaged bits of the representation, etc.

[2] Block (2023, p. 225) notes experimental evidence against models of temporal analogue magnitude representation that *would* satisfy the Parts Principle. See Block (2023, p. 223) and Burge (2018) for further issues with this conception of iconicity.

Block favors instead a conception of iconicity derived from Roger Shepard and colleagues' emphasis on 2nd-order isomorphism (Shepard & Chipman 1970; Shepard 1978). Central to this conception is that relations among representations themselves mirror relations among the particulars or features that the representations represent. For example, dots on a subway map are related spatially in a way that mirrors spatial relations among the stations the dots represent. Likewise, the relations among a thermometer's varying mercury heights (higher than, lower than) mirror the relations among the varying ambient temperatures that they represent (hotter than, colder than). Again, linguistic representations do not generally work this way: the words 'red' and 'green' do not mirror relevant relations between the colors.

Here's Block's full statement of his "analog tracking and mirroring" conception of iconicity:

> Analog tracking and mirroring obtains when there is a set of environmental properties and a set of representations of those environmental properties such that:
>
> 1. Certain differences in representations function as responses to differences in environmental properties in a way that is sensitive to the degree of environmental differences. …
>
> 2. Certain differences in representations function to alter the situation that is represented in a way that depends on the degree of representational change.
>
> 3. Certain relations (including temporal relations) among the environmental properties are mirrored by representations that instantiate analogs of those relations. (Block 2023, pp. 221-2)

Crucially for our discussion, iconic representations on Block's view represent properties that come in *degrees*, and they do so by coming in families of representations themselves organized by *degree* of vehicular change. Degree of change in representational content corresponds (as a matter of function) to degree of change in representational vehicle and vice versa. Indeed, Block emphasizes that "[t]he key difference [between iconic and discursive representations] comes in with the role of *degrees of difference*" (Block 2023, p. 222—italics in original). Iconic representation, unlike non-iconic representation, thus requires a family of properties organized along some graded dimension(s) and a family of representational vehicles likewise arrayed along some graded dimension(s) such that the placement of features and vehicles, respectively, along these dimensions relevantly mirror one another.

Peacocke's counter-example to the Parts Principle provides an illustration: The loudness of sounds comes in degrees, and iconic representations as of loudness in principle might be realized in a correspondingly graded way by average number of spikes by some neural population across a temporal interval: the more spikes, the louder the sound is represented to be. Different spike rates would correspond to different volumes, and the representations would be organized along a dimension—number of spikes per unit time—that mirrors the greater or lower volumes of the sounds.[3]

[3] Peacocke's example shows that iconicity in Block's sense does not entail satisfaction of the Parts Principle. Does satisfaction of the Parts Principle entail iconicity in Block's sense? *Modulo* details of formulation, it would seem that

What is meant here by coming in *degrees*? Block leaves the notion undeveloped, beyond noting that density is not required and magnitudes can be "digitized" (his quotes—Block 2023, p. 222). Rather than attempt an explication on his behalf, I'll hope that any mistaken understanding on my part will occasion useful clarification. But I will take seriously the idea that the relevant environmental properties are in some sense degreed and that an accurate perceptual attribution of these properties must in some sense capture this and do so in part in virtue of degreed aspects of the representational vehicle, as per clauses (1) and (2) of Block's characterization of iconicity. (As Block (2023, p. 222) notes, clause (3) also involves degrees: the "mirroring of degrees by degrees".) In section 4, I'll consider a weakening.

3. Perceptual Categorization

The *prima facie* problem is that there seem to be perceptual attributions that are not iconic in Block's sense. For example, Block, in agreement with many others, maintains that we perceptually attribute being a face, being a certain phoneme, being an object (or being a Spelke-object), being inside of, being the cause of, etc. But, in each case, it's not obvious that the requirement of gradedness is satisfied. I focus on perceptual categorization.

Perceptual categorizations are perceptual attributions as of belonging to a certain category. The term 'category' gets explicated in various ways in these contexts. For example, Burge (2022, p. 487) limits perceptual categorization to attribution of kinds, while Block (2023, p. 271) includes color categorizations—attributions of features, not kinds, on Burge's view. What matters for us is just whether the attributions are iconic in Block's sense, however they are themselves categorized. Again, a reason for thinking they are not iconic is that they seem to attribute non-graded attributes. Some actually build non-gradedness into their characterization of categorization—e.g., Mather (2016, p. 143): "Each stimulus is perceived as a member of a discrete category, in an all or nothing fashion, rather than occupying a position along a continuous dimension." But it's more informative to explore arguments for the claim, rather than assume it from the start.

A first argument is that certain ways of speaking or thinking that would indicate gradedness seem unacceptable. For example, it seems unacceptable to speak or think of one thing as more of a face, or "facier", than another.[4] Such considerations are indeed a common source of evidence in linguistic semantics—for example, in discussions of graded adjectives (Kennedy and McNally 2005). But, as this very fact reminds us, they fall short of establishing ungradedness in perceptual attribution. The unacceptability judgments may just reflect features of our corresponding concepts and/or the lexical items that express them. That a conceptual attribution, or its linguistic

clause (3) would be satisfied insofar as the mereological structure of the representations mirrors the mereological structure of what they represent—supposing mereological relations count as among clause (3)'s "certain relations" and would suffice. (Thanks here to Jeremy Goodman.) It is less clear that satisfaction of the Parts Principle entails satisfaction of clauses (1) and (2)—in part because of the degrees requirement again. *Prima facie*, a representation may satisfy the Parts Principle while representing something as having some non-graded property.

[4] To be sure, one face may be bigger than another or closer to a statistically normal or paradigmatic face than another. But—the claim would be—these are distinct from being more of a face in the sense intended. One might also argue that we perceptually attribute, not facehood, but something like face-shaped or face-like, which may well come in degrees. (Such attributions would be accurate of smiley faces—perhaps to some degree.) But this seems not to be Block's view (Block 2023, p. 82f.).

expression, does not attribute a graded property does not entail the same of a corresponding perceptual attribution.

What could follow—depending on how Block might develop his talk of degrees—is that what's attributed in perception is not the same as what's attributed conceptually in such cases. This would be so, for example, if degreed attributions attribute degreed properties, but non-degreed attributions don't: one and the same property can't be both graded and ungraded. Or if, say, what degreed attributions attribute are relations between a particular and position on a scale, while what non-degreed attributions attribute are monadic. The transition from the perceptual attribution to the corresponding conceptual attribution—a transition perhaps mediated by a graded cognitive attribution—would then involve in that respect a change of content. Moreover, when graded perceptual attributions were characterized using ungraded language (for example, perceptual attribution as of being a face, *sans* degree), we would want to understand this as a *façon de parler* that exploits a ready-to-hand lexical expression to indirectly indicate its corresponding perceptual attributive.

A second way to argue for non-graded perceptual attribution adverts to empirical demonstrations of categorical perception in the technical sense. Categorical perception in the technical sense is when equidistant items or features along a certain objective dimension are perceived as more similar within certain boundaries than across them. These boundaries are then theorized to be the boundaries of categories. For example, one can vary relevant aspects of the acoustic signal continuously, but have subjects judge language-like sounds within boundaries to be more similar than objectively equidistant sounds across boundaries. In part on this basis, it's claimed that subjects perceive phonemes in a way that conforms to there being fairly sharp boundaries between them.

But does categorical perception in fact challenge Block's claim that perceptions are constitutively iconic in his sense? Categorical perception is often described as a *warping* of the represented feature space relative to how the features are related objectively: some parts of the space are stretched, others smushed, relative to the objective space (Goldstone and Hendrickson 2010; Kronrod et al. 2016). For example, color hue may vary continuously, and our perceptual representations of color may be graded (down to the limits of acuity); but, if there is categorical perception of color (an unsettled question—cf. Witzel 2019, also McMurray 2022 on phonemes),[5] the "distance" between color hues in the represented space may be deformed in a way that leads to the judgments characteristic of categorical perception. If so, further argument, beyond the existence of categorical perception in this sense, would be needed to conclude that perceptual representation of colors isn't graded. Moreover, it's worth noting that the deformations in this case would preserve *ordinal* relations among colors, satisfying clause (3)—the "mirroring" clause—in Block's characterization of iconicity.

[5] McMurray (2022) argues that the patterns of judgment characteristic of categorical perception are highly task-dependent in the phonemic case and that the relevant perceptual encodings are not warped. Denying categorical perception is of course another strategy in reply to an argument that invokes it—albeit one not available to Block without alteration of his *own* arguments invoking it (Chapter 6). McMurray, though he challenges categorical perception of phonemes, affirms phonemic categorization and allows that it might be perceptual. He also argues, however, that phonemic categorization is graded.

One further argument would advert to downstream effects. If representational states between "boundaries" share their downstream effects, that would provide *prima facie* reason to assign them the same content—for example, *red* as opposed to, or in addition to, more fine-grained contents corresponding to degrees of redness. Similarly, one may argue that what matters crucially downstream in linguistic comprehension (in particular, for lexical identification) is what phoneme is represented—not, or not just, the graded lower-level features in the acoustic signal on which phoneme perception in part depends in an extremely complex way (Kazanina et al. 2018).

This is not the place to delve into the empirical details necessary to develop and evaluate such arguments. But I'll mention two strategies in reply. The main point is just to indicate—what's perhaps obvious—that these matters involve substantial, unobvious empirical commitments.

One reply would argue that these downstream effects and, importantly, the representations responsible for them are post-perceptual. This would seem to be in tension with other of Block's commitments: Block (2023, pp. 64-9—and cf. pp. 271-3) discusses at length the susceptibility of phonemic representations to adaptation, evidence that they are perceptually attributed. But perhaps the tension could be resolved by positing degreed perceptual phonemic attributions (susceptible to adaptation) that transition to non-graded post-perceptual phonemic attributions that are more directly responsible for the downstream effects—similarly for other cases. To be compelling, this obviously would require empirical backing. Alternatively, one might try reinterpreting the phonemic adaptation results in terms of adaptation to lower-level features (cf. Block 2023, p. 65).

The other strategy would argue that these downstream effects can result from computations over graded representations. There is indeed significant work on how to computationally relate graded and non-graded representations (e.g., Smolensky et al. 2014). Here, though, I would caution that not all such work sees graded representation as thereby attributing graded properties or otherwise including gradedness in the representation's content, as Block's clause (2) would seem to require. That a representation of /b/, for example, can be activated to varying degrees need not entail that an instance represents *being /b/ to such-and-such degree* (Kazanina et al. 2018). Among other possibilities, the level of activation might correspond to a perceptual credence (Feldman et al. 2009).[6]

Dropping Talk of Degrees

We've been exploring ways of arguing that perceptual categorizations are not graded, as well as responses thereto. The discussion has put much weight on Block's talk of degrees, appropriately so given his own emphasis. But there's another line of reply worth exploring—one that just *drops* this talk of degrees. This might count as a *reply*, not a capitulation, if it would otherwise preserve what's crucial in Block's characterization of iconicity. (It would definitely count as a reply, not a capitulation, if it in fact captures what Block intends by his admittedly undeveloped talk of degrees, while just dispensing with a possibly misleading term.)

[6] Block (2023, pp. 200-14) argues for the viability of instrumentalism regarding Bayesian inference in perception, in part in response to Gross (2020). Some of his discussion suggests he would extend this to probabilistic perception more generally. In any event, Block's (2018) competition model suggests gradiently-activated representations without representation of degreed properties.

The idea is that iconicity might require just that there be a family of representations arrayed along representational dimensions that mirror some environmental dimension(s)—where the representations function to be sensitive to differences along the environmental dimension(s), and changes of representational vehicle function to change representational content along the relevant dimension(s). Iconicity might require *that*, without requiring further that the representational changes and environmental differences come in degrees. Phonemes, for example, could find their place in a phonemic space. But it needn't be that one sound is more of a /t/ than another or that perceptual phonemic attributions represent or entail that they are—even if the sounds may come in degrees along the *lower*-level dimensions that constitute the space. (Compare the overlay of political boundaries on a spatial region: locations may be closer and further from one another, but either in a state or not, not to some degree.) Phonemic representations could then be iconic insofar as the representations form a space that mirror relations among the phonemes and represent the phonemes in part in virtue of doing so—without requiring phonemes, or representations thereof, to be graded. This line of reply would have the added benefit of handling not just *prima facie* non-graded perceptual categorizations, but also other apparent non-graded perceptual attributions.

But there's a new problem. There seem to be perceptual attributions that don't (or don't obviously) fall into a *family* of such representations—representational loners, if you will. Examples might be highly generic categories or features, such as being colored or shaped. Representations as of being red fall into a family with other color representations. But what of a representation as of being colored—not as of being this or that color, but simply as of being colored? (Similarly, for a representation as of being shaped.) What is the family to which it belongs, such that members of the family are arrayed in a space and it is empirically plausible that they represent similarly arrayed properties in part in virtue of falling into such a family with vehicles thus arrayed? Neither the family of *most generic perceptual determinables* nor representations thereof seem so arrayed.

Of course, such cases pose no problem if in fact we do not make such highly generic perceptual attributions. Block may deny that we perceptually attribute being colored or being shaped—*contra* Burge (2022). But there are perceptual representations Block does defend which likewise may be loners. An example may be perceptual attributions as of something's being an object, or as of being a Spelke-object: a relatively rigid, cohesive 3-d body with a closed boundary (Spelke 1990). What is the family to which perceptual attribution of objecthood belongs, such that members of the family are arrayed in a space and it is empirically plausible that they represent properties similarly arrayed in part in virtue of falling into such a family with vehicles thus arrayed?

It might be tempting to reply that loners are just *limit* cases: a family of one kind of representation, related to itself and perhaps its absence.[7] As such, they might seem to satisfy our weakened clauses for Block's conception of iconicity. Clauses (1) and (2) are satisfied insofar as a change in Spelke-objecthood tends to lead to a change in perceptual representation, and a relevant change in the perceptual representation changes whether it represents something as a

[7] Absence: not representing something as an object should not be confused with representing something as not an object. Block (2023, pp. 182-6) denies that there is negation in perceptual content.

Spelke-object. As for clause (3), perhaps Spelke-objects' being of the same kind as one another is mirrored by representations thereof likewise being of the same kind. The problem, however, is that this trivializes iconicity so that even paradigmatically discursive representations count as iconic. Recall Block's underscoring how gradedness is key to distinguishing iconic and discursive representations. On our weakening, multiplicity is needed to play the same role.

Here's a different reply to the *prima facie* problem of loners. Attributions of Spelke-objecthood are higher-level in that they depend on the attribution of *other* properties: cohesiveness, rigidity, and boundedness. Perhaps the iconicity of the lower-level attributions on which such higher-level attributions depend might suffice to render the *higher*-level attributions iconic. What would have to be the case for this to be so? Suppose entities can be arrayed along various dimensions according to their cohesiveness, rigidity, and boundedness. Spelke-objects would be limit cases, or at least cases sufficiently cohesive, rigid, and bounded. Now consider similarly-arrayed *representations* of these entities. These representations form a family related in ways that mirror relations among the properties of the entities represented. Suppose further that representations as of cohesiveness, rigidity, and boundedness are iconic. (This might be unobvious—perhaps one would need to consider even lower-level attributions. But bracket this.) The question is whether that would suffice to render attributions of objecthood themselves iconic.[8]

Not yet. For the attribution of objecthood might involve a representation (a vehicle) distinct from the representation of the lower-level features. Having represented an entity as cohesive, as rigid, and as bounded, perception may then transit to a distinct state that represents it as an object. There's no reason to think the iconicity of the lower-level representations renders the distinct higher-level representation iconic.

But what if we suppose further—as an empirical hypothesis—that the state that represents objecthood is the very state that represents the lower-level features? The idea is not to reduce the representation of objecthood to that of the lower-level features. That would be a contentious claim—and we can readily conceive of organisms that represent the lower-level features without representing the higher-level feature. (Cf. Block's (2023, p. 96) point that "recognitionally equivalent" congeries of low-level attributions do not necessarily eliminate high-level attributions.) Rather, the suggestion is that each state that represents the high-level property is token-identical to a state that represents the relevant lower-level properties. It may also have the higher-level content in virtue of the state's functional role—for example, downstream effects on, say, tracking or grasping. (As mentioned, we can imagine an organism that represents the lower-level features, but not objecthood: the downstream effects associated with objecthood would then be absent.) A system needn't be built that way: again, perception could instead transit to a *distinct* state. But suppose it were. The state that represents objecthood would then be iconic, since it is being granted that the state iconically represents the relevant lower-level features.

This reply rests on a lot of 'supposes' (a further one is that all *prima facie* loners are higher-level). But, in addition, it has an apparently objectionable feature. Representations of objecthood, on this reply, would be iconic in virtue of the iconicity of *other* aspects of the representations that represent objecthood—albeit aspects on which representation of objecthood

[8] Similarly for other cases, such as perceptual attribution of facehood and the fifty lower-level dimensions it is hypothesized that this involves—Tsao 2019.

in part depend. That is, the reply does not show how to locate objecthood and representations thereof within an appropriate family. The claim is rather that, as it happens, representations as of objecthood are always representations that are iconic *in other ways*. Consider a labeled map—perhaps a weather map containing 'HUMID' by a dot labeled 'Houston'. One wouldn't consider the attribution of humidity and the identification of Houston as such iconic just because the map has other iconic aspects. The map's representational format is hybrid, with both iconic and non-iconic aspects.[9]

A final reply is to claim—or, better, to argue, as Block does—that Spelke-objecthood and representations thereof in fact *are* graded, so that the relevant families are comprised of differing degrees of objecthood and correspondingly degreed representations.[10] This might seem a return to the position of the previous section, rather than a distinct reply based on a weakening of Block's iconicity clauses. But it needn't be: the reply can be localized to perceptual loners. That is, one might admit ungraded perceptual attributions and deploy the weakened clauses to accommodate them, and also argue that all perceptual "loners" *are* graded—and so not really loners. (One could combine this with denying that some *prima facie* cases are *perceptual*—as perhaps with the highly generic attributions.) This reply would be strengthened to the extent one could provide a rationale for why all "loners" should be graded, as opposed to resting on its happening to be the case that they are.

To conclude: I have provided some development of the worry that Block's conception of iconicity does not apply to cases he'd consider perceptual. This is a worry because Block maintains that perception is constitutively iconic. I have also explored a variety of replies and their potential costs and burdens, hopefully in a way that fruitfully furthers development of the view. But it could be that none of them appeal. What do you think, Ned?[11]

References

Beck, J. 2023. Review of Ned Block, *The Border between Seeing and Thinking*. *Notre Dame Philosophical Reviews.*

Block, N. 2018. If perception is probabilistic, why doesn't it seem probabilistic? *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*. https://doi.org.10.1098/rstb.2017.0341

Block, N. 2023. *The Border between Seeing and Thinking*. Oxford: Oxford University Press.

Block, N. forthcoming. Word Representations and Ampliative Transitions: Reply to Steven Gross. *Analysis Reviews*.

[9] There's an objection to my analogy, but Block (2023, p. 217) in effect anticipates and supplies a response:

> … perhaps a map with labels is really two distinct representations, an iconic representation that represents, e.g., rivers and their shapes, and a set of discursive labels. [Ian] Phillips noted, however, that the shape of a river could be denoted on a map by a shaped string of letters spelling out the river's name.

The 'two representations' suggestion would need also to account for role played by *placement* of the labels.

[10] Block (2023, p. 242) adverts to Marino & Scholl's (2005) result that contour-based lower-level cues can drive object-like attention even in the absence of objecthood, but don't fully account for object-based attention. As my characterization suggests, it is unobvious that the results support graded objecthood.

[11] Thanks to Jake Beck, Justin Bledin, Jeremy Goodman, EJ Green, Jake Quilty-Dunn, Ian Phillips, Kyle Rawlins, and Susanna Schellenberg. Special thanks to Ned Block for conversation, encouragement, and his fantastic book, about which one could easily write multiple commentaries without exhausting its riches. And one has—see Gross (forthcoming) and Block (forthcoming) in reply.

Burge, T. 2018. Iconic representation: Maps, pictures, and perception. In Wuppuluri Shyam & Francisco Antonio Dorio (eds.), *The Map and the Territory: Exploring the Foundations of Science, Thought and Reality*. Springer. pp. 79-100.
Burge, T. 2022. *Perception: First Form of Mind*. Oxford: Oxford University Press.
Carey, S. 2009. *The Origins of Concepts*. Oxford: Oxford University Press.
Feldman, N., Griffiths, T., and Morgan, J. 2009. The influence of categories on perception: Explaining the perceptual magnet effect as optimal statistical inference. *Psychological Review*. 116: 752-82.
Firestone, C. and Phillips, I. 2023. Seeing fast and thinking slow: A philosopher explores perception and cognition. *Science* 379: 1196.
Fodor, J. 2007. The revenge of the given. In B. McLaughlin and J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind* (Oxford: Blackwell), 105–116.
Goldstone, R. and Hendrickson, A. 2010. Categorical perception. *WIRES Cognitive Science*. 1: 69-78.
Green, E. J. 2023. "The Perception-Cognition Border: Architecture or Format?" In B. P. McLaughlin & J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*, 2nd ed. Oxford: Blackwell, 469-93.
Green, E.J. and Quilty-Dunn, J. 2017. What is an object file? *British Journal for the Philosophy of Science* 72: 665-99.
Gross, S. 2020. Probabilistic Representations in Perception: Are There Any, and What Would They Be? *Mind & Language* 35: 377-89.
Gross, S. forthcoming. Language and the border between perception and cognition. *Analysis Reviews*.
Kazanina, N., Bowers, J., and Idsardi, W. 2018. Phonemes: Lexical access and beyond. *Psychological Bulletin & Review* 25: 260-85.
Kennedy, C. and McNally, L. 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language* 81: 345–381.
Kosslyn, S. 1980. *Image and Mind*. Cambridge: Harvard University Press.
Kronrod, Y., Coppess, E., and Feldman, N. 2016. A unified account of categorical effects in phonetic perception. *Psychonomic Bulletin & Review*. 23: 1681-1712.
Marino, A. and Scholl, B. 2005. The role of closure in defining the "objects" of object-based attention. *Perception & Psychophysics.* 67: 1140-9.
Mather, G. 2016. *Foundations of Sensation and Perception*, 3rd Edition. New York: Routledge.
McMurray, B. 2022. The myth of categorical perception. *Journal of the American Acoustic Society.* 152: 3819-42.
Peacocke, C. 2019. *The Primacy of Metaphysics*. Oxford: Oxford University Press.
Shepard, R. 1978. The mental image. *American Psychologist* 33: 125-137.
Shepard, R, and Chipman, S. 1970. Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology* 1: 1-17.
Smolensky, P., Goldrick, M., and Mathis, D. 2014. Optimization and Quantization in Gradient Symbol Systems: A Framework for Integrating the Continuous and the Discrete in Cognition. *Cognitive Science* 38: 1002-38.
Spelke, E. 1990. Principles of object perception. *Cognitive Science* 14: 29-56.
Tsao, D. 2019. Face values. *Scientific American*. 2029: 23-9.
Witzel, C. 2019. Misconceptions about color categories. *Review of Philosophy and Psychology* 10: 499-540.