Perceptual Abstraction

Mason Westfall & E. J. Green
Penultimate Draft (forthcoming, *Noûs*)

Abstract: Perception puts us in touch with highly determinate properties of objects, such as fine-grained color shades and detailed surface shapes. However, most of our immediate perceptual judgments concern more abstract properties, such as the property of being a dog, or of being red *simpliciter*. So, how does perception attune us to abstract properties despite its evident determinacy? This paper argues that perception can be sensitive to abstract properties in multiple, fundamentally different ways. We articulate a distinction between *implicit* and *explicit* perceptual abstraction and explore its ramifications for debates about the role of concepts in perception, the richness of perceptual content, and rationalist versus empiricist views of our innate representational repertoire. We also propose an empirical test for differentiating implicit from explicit perceptual abstraction. Finally, we outline a version of Conceptualism built on the implicit/explicit distinction, on which explicit perceptual abstraction involves the use of concepts, while implicit abstraction can occur non-conceptually. We argue that this form of Conceptualism is both empirically viable and unscathed by prominent Non-Conceptualist arguments.

1. Introduction: The Generality of Perception

Perception puts us in a position to think about abstract properties, such as kinds and determinables. When you see a McIntosh apple, you are in a position to form beliefs not just about its specific shape or shade of red, but about its property of being an apple, or its property of being red *simpliciter*. At the same time, perception seems "determinate", always presenting us with specific shades of red, and never just redness *simpliciter*. How does perception enable us to rapidly cognize abstract properties, despite its evident determinacy? In our view, perception prepares us to cognize abstract properties in multiple, fundamentally different ways. An appreciation of the varieties of perceptual abstraction has implications for debates about the richness of perceptual content, the role of concepts in perception, and rationalist and empiricist accounts of concept acquisition.

According to a venerable tradition, abstraction is a post-perceptual achievement. Introspection may seem to motivate this idea. After all, while we enjoy experiences of determinate shades of red, we never seem to experience *just* the determinable redness. Nothing ever looks red without looking some specific shade of red. So, perhaps the most reasonable hypothesis is that experience only represents determinate colors, not determinable color categories. As Michael Thau puts it:

> [I]t's completely obvious that, even if . . . colors are represented in visual perception, the color red . . . is never represented in visual perception; rather, objects are visually represented as being some determinate shade of red (Thau 2002: 217)

For some, perception's specificity and cognition's relative abstractness suggests a deep difference between them. Our ability to perceptually represent fine-grained properties for which we lack concepts might indicate that perception—unlike cognition—represents information non-conceptually (Evans 1982; Heck 2000). Or perhaps it indicates that perception represents information in an iconic format, as contrasted with cognition's discursive or digital format. As Dretske puts it:

> Perception is a process by means of which information is delivered within a richer matrix of information (hence in *analog* form) to the cognitive centers for their selective use. Seeing, hearing, and smelling are different ways we have of getting information…to a digital-conversion unit whose function it is to extract pertinent information from the sensory representation for purposes of modifying output… The traditional idea that knowledge, belief, and thought involve concepts while sensation (or sensory experience) does not is reflected in this coding difference. (Dretske 1981: 30)

The idea that perceptual representation is highly specific also forms the backdrop for empiricist theorizing about concept learning. For empiricists, the challenge is to explain how we come to have general representations at all, such as the concept of redness, given only a diet of perceptual encounters with determinate shades of red. Locke's answer was that we must *abstract* the concept of redness from perceptual experiences of red things. On this view, perception itself does not present us with abstract properties like redness *simpliciter*. Further processing is needed to set abstract properties before the mind.[1]

However, even if our perceptual states always represent highly determinate properties, it doesn't follow that they *only* represent such properties. To the contrary, both introspection and empirical investigation have been taken to indicate that abstract properties can be apprehended perceptually. Consider color categories. Many observers report that a rainbow appears to contain discrete "stripes" with relatively clear boundaries, though the light reaching the eyes varies smoothly in wavelength from the top to the bottom of the rainbow. Additionally, empirical studies of *categorical perception* reveal that we are sometimes better at distinguishing stimuli that span categorical boundaries than stimuli that fall within the same category, even when "objective" similarity is controlled (Liberman et al. 1957; Harnad 1987; Goldstone and Hendrickson 2010). Gradual changes to the physical properties of an auditory stimulus change the heard syllable from /be/ to /de/ to /ge/. While the physical variation is gradual, allegedly the perceptual variation is not. Rather, perception discounts differences within the same phonemic category while exaggerating differences across categories (Liberman et al. 1957; Swan & Myers 2013; although see McMurray 2022). Such results support the view that abstract categories at least make a perceptual difference. But what kind of perceptual difference do they make?

[1] See Locke (1690/1975: II, xii, §1), Mackie (1976: ch. 4), and Gauker (2011: ch. 1).

This paper distinguishes multiple ways that perception can be sensitive to abstract categories and explores the implications of these distinctions. Section 2 introduces a distinction between implicit abstraction and explicit abstraction: Implicit perceptual abstraction occurs when perceptual representations of determinate properties are organized in ways that produce differential sensitivity to abstract categories, while explicit abstraction occurs when perception forms an explicit representation of a category. Section 3 draws out consequences of the implicit–explicit distinction for debates about the role of concepts in perception and the makeup of our innate representational repertoire. Section 4 proposes an evidential criterion, which we call the *Dissimilarity Test*, for differentiating implicit from explicit abstraction, and applies the criterion to the perception of abstract shape categories. Section 5 leverages the implicit–explicit distinction to articulate a version of Conceptualism on which explicit perceptual abstraction involves the use of concepts, while implicit abstraction can occur non-conceptually. Section 6 concludes.

2. Implicit versus Explicit Abstraction

By *perceptual abstraction*, we'll mean, very roughly, any case in which perception exhibits differential sensitivity to a property of an object that is more abstract than certain other properties that the object is also perceived as having. For example, if we perceptually represent a given surface as both navy blue and blue *simpliciter*, then we exhibit perceptual abstraction with respect to color. This section explores multiple ways that perception can display differential sensitivity to abstract properties.

We construe perceptual abstraction as a *state* or *achievement*, not as a *process*. Many theorists use the term "abstraction" to refer, roughly, to a process of acquiring representations of more abstract properties on the basis of representations of less abstract properties (e.g., Mackie 1976: ch. 4; Laurence & Margolis 2012; Sundström 2019). We do not require that abstract properties be learned or even computed from more determinate properties. For example, suppose that the perceptual representation of abstract shape properties (e.g., *triangle*) is computationally independent of the representation of determinate shape properties (e.g., *equilateral triangle*).[2] A process-oriented conception might not count this as abstraction, since no process takes a more determinate representation as input and outputs a less determinate representation. For us, a perceptual system that represents both determinate and abstract shape properties exhibits "abstraction", even if abstract properties are computed independently of more determinate properties, and even if the ability to perceptually represent abstract shape properties is innate.

There are at least two ways in which one property can be *more abstract* than another. First, properties within a single dimension (or multi-dimensional feature space) that stand in determinable–determinate relations (e.g., *red* and *scarlet*) are classified as more or less abstract. In these cases, the determinates and the determinables carve out points (or regions) of the same

[2] In fact, some have suggested that the visual system operates in *reverse* order, recovering more abstract shape properties prior to computing more determinate properties (see also Chen 2005).

dimension (or multi-dimensional space), with the determinates carving out points (or regions) wholly contained within those carved out by the determinables. Second, kinds such as *dog* or *water* are standardly classified as more abstract than the lower-level features, such as color, shape, or texture, which probabilistically signal membership in those kinds. In both cases, more abstract properties are preserved despite differences in their less abstract counterparts. Burgundy and scarlet are both shades of red, despite fine-grained differences. And Beagles and German Shepherds are both dogs, despite differing along many of the lower-level dimensions that we rely on to classify them as dogs.

Perception can facilitate cognition of abstract properties either by explicitly representing them, or by organizing representations of determinate properties in ways that facilitate their explicit representation, while leaving the formation of these explicit representations up to post-perceptual cognition.

The prototypical way to explicitly represent a property is to form a representation that *extracts* the property. Following Dretske (1981: ch. 7), we'll say that a representation R extracts a property F when either R, or a separable constituent of R, represents F without representing anything else, or anything more specific (see also Matthen 2005: ch. 3; Kulvicki 2007). For example, the sentence "the red house is to the left of the brown shed" has separable constituents that represent the properties *red* and *brown*, and nothing more specific. These constituents are "separable" because they can remain invariant while other parts of the sentence are dropped or modified. Conversely, a richly detailed photograph of a red house may in some sense encode the house's redness, but it has no separable constituent that represents redness *simpliciter*. The very same parts of the photograph that encode the house's redness also encode information about its determinate shade of red. When a representation extracts an abstract property, we'll call it a *category label*.

In implicit abstraction, perception is attuned to abstract categories without explicitly representing them. This attunement makes it easier for downstream processes to determine whether a stimulus belongs to a given category. We distinguish two forms of implicit perceptual abstraction: *content*-based category attunement and *vehicle*-based category attunement.

2.1 Content-Based Category Attunement

Many of the properties we perceptually represent fall into families—colors, shapes, sizes, orientations, etc. Properties within a family *exclude* one another. If a surface is square, then it can't be circular; and if it is matte, then it can't be glossy. Conversely, properties from distinct families normally do not exclude one another. A surface's shape does not constrain the color it can possess. When we speak of *dimensions*, we will have in mind families of this sort.[3]

[3] This is a rough characterization, not a definition. Two properties may exclude one another without falling along the same dimension. The property of being a triangle excludes the property of being a cat, even though *triangle* and *cat* do not fall along a single perceptible dimension (see Green 2020: 327).

Many perceptible dimensions are *quantitative*, insofar as degrees of dissimilarity between values of the dimension can be numerically measured and compared. For example, determinate colors can be understood as points in a three-dimensional color space corresponding to hue, brightness, and saturation, where distances between points in this space correspond to the properties' perceptual dissimilarity (Clark 1993: ch. 4; Rosenthal 2010; Gauker 2011: ch. 5). Colors further apart in the space are more dissimilar than those closer together. This captures the sense in which purple is more similar to red than it is to yellow.

Content-based category attunement involves a perceptual system exhibiting biases in the representation of a dimension (or a multidimensional similarity space) that facilitate the extraction of some abstract category. Perception might be biased toward representing certain values of a dimension over others, and these biases might function to enhance perceptual dissimilarities between objects belonging to different categories relative to those belonging to the same category. Content-based attunement is so called because it involves categorical sensitivity at the level of perceptual *content*.

One form of content-based attunement involves selectively *stretching* perceptible dimensions that are relevant to category membership, accentuating perceptible differences along these dimensions relative to differences along dimensions that are category-irrelevant (Goldstone 1994; van Gulick & Gauthier 2014; Folstein et al. 2015). For example, if shape is more important than other dimensions for discriminating among vehicle categories (sedan, SUV, etc.), then shape differences among vehicles may be perceptually exaggerated, effectively stretching the shape dimension for that class of stimuli, making the stimuli appear more dissimilar in shape than they really are.[4]

Another form of content-based attunement, called *perceptual magnetism* (e.g., Kuhl 1991; Feldman et al. 2009), involves a systematic bias against representing values of a dimension that are near category boundaries and toward representing values that are near category centers or "perceptual magnets." For instance, an uttered syllable that is *in fact* very near the /be/-/de/ boundary might be perceptually misrepresented as having a value closer to the center of the /be/ category. Consequently, the perceptual space of phonemes would be compressed near category centers and expanded near category boundaries, relative to objective differences among the stimuli.

Perceptual magnetism can be implemented in multiple ways (see Bates & Jacobs 2020; Dubova & Goldstone 2021). One way, inspired by Bayesian models of perception, involves a non-uniform perceptual prior over the relevant dimension (or multidimensional space) that systematically assigns higher probability to reference points at or near category centers, falling off toward category boundaries (Feldman et al. 2009). The prior would bias the perceptual system toward representing objects as closer to category centers than they really are. As a result, perceptual similarity space would effectively "shrink" near the centers of categories, making stimuli near category centers

[4] For behavioral evidence supporting dimension-specific stretching, see Goldstone (1994), Goldstone & Steyvers (2001) and Folstein et al. (2012); for neuroimaging evidence, see Folstein et al. (2013). See also Connolly (2019: 19-28), Burnston (2023), Jenkin (2023), and Ransom and Goldstone (2024) for helpful discussion.

appear more similar than they really are. Conversely, stimuli on opposite sides of category boundaries would be "pulled" in opposite directions, toward different category centers, making them appear more different than they really are (see figure 1).

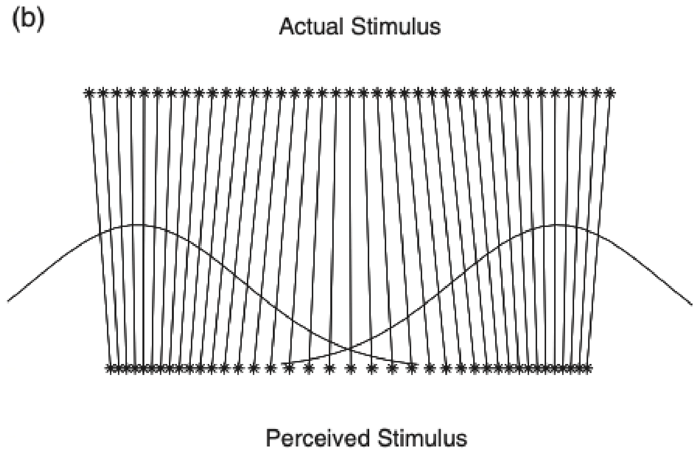

Figure 1. Illustration of how a non-uniform prior peaked around category centers produces a shrinking of perceptual similarity space near category centers and an expansion of perceptual similarity space near category boundaries. *Source*: Feldman et al. (2009). Material in the public domain.

Some studies have uncovered dimension-specific stretching and perceptual magnetism operating in tandem. For example, Dubova & Goldstone (2021) trained participants on a categorization task involving a morph series of fish, where images in the series varied along two dimensions, only one of which was relevant to categorization. When a fish was presented, the participant was asked to reconstruct it by adjusting a second fish image. The results indicated that perceptual representations were both more precise along the category-relevant dimension than the category-irrelevant dimension (as predicted by dimension-specific stretching) and assimilated toward category centers for the category-relevant dimension (as predicted by perceptual magnetism).

Content-based attunement can produce perceptual sensitivity to a category without perceptual extraction of that category. To illustrate, consider a recent study due to van Geert and Wagemans (2023). Van Geert and Wagemans made a series of morphs between paradigms of shape categories (see figure 2). For example, one morph series interpolated between a paradigmatic car shape and a paradigmatic turtle shape, with physical differences held constant between adjacent morphs. Subjects were shown pairs of morphs and asked for either a same/different judgment or a similarity rating. Both discrimination accuracy and similarity ratings were significantly influenced by the distance of the morphs from the paradigms. Pairs of stimuli close to paradigms were more difficult to discriminate, and rated as more similar, than pairs of stimuli further away from paradigms. A derivative "category boundary effect" was produced, whereby stimuli crossing category boundaries are easiest to discriminate, because morphs straddling the boundary are furthest from paradigms. However, the effect is a linear function of distance from the paradigms, and operates even for pairs of within-category comparisons. Specifically, subjects were better at discriminating within-category pairs farther from category paradigms than those closer to paradigms. This suggests that perceptual sensitivity to the contrast between within- and between-category pairs emerges from perceptual magnetism toward category prototypes, rather than from a binary distinction between categories.

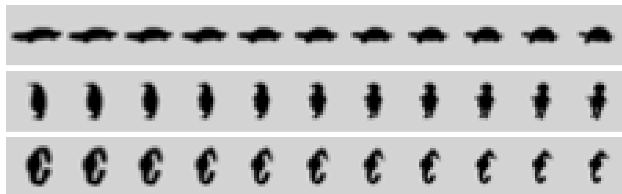

Figure 2. Morph series used by van Geert and Wagemans (2023). Reprinted with permission of the American Psychological Association.

Perceptual magnetism can explain these results, even without explicit representation of the relevant shape categories. We need only assume that perceptual similarity increases (and discriminability decreases) as stimuli get closer to reference points or category paradigms.[5] As illustrated in figure 1, this pattern would be predicted by the hypothesis that the perceptual system imposes a non-uniform prior on the underlying shape space, peaked at reference points. We needn't further assume that perception forms representations that extract the relevant shape categories (e.g., *car shape* or *turtle shape*).

This is not to say that perceptual magnetism is *inconsistent* with category extraction. It's possible that perception extracts categories *and* imposes a non-uniform prior peaked at category centers. Alternatively, biases toward category centers might be produced through other mechanisms involving category extraction. For example, Bae et al. (2015) found a bias toward representing color shades near category centers, both in perception and in visual working memory. The authors advocate a "dual content model", on which color estimates are the result of integrating explicit color category representations with specific values or "point estimates." Because this model appeals to explicit color category representations, it explains perceptual magnetism partly in terms of explicit abstraction. However, as Bae et al. also note (18), one could alternatively explain their results solely via a non-uniform prior over determinate shades.[6]

Content-based category attunement facilitates explicit abstraction at post-perceptual stages for at least two reasons. First, if between-category differences are exaggerated and within-category differences are downplayed, then between-category discrimination will be relatively easy and within-category discrimination relatively hard (even without explicitly representing categories).

[5] Because discrimination and similarity judgments are potentially subject to extra-perceptual influences, it is possible that these biases toward category paradigms emerged only at the level of cognition, not within perception. However, our present point is just that *if* the relevant patterns of categorical sensitivity are perceptual, then they can be explained by perceptual magnetism without appeal to explicit category representation.

[6] Cp. Green (2024).

Second, if perceptual representations are biased toward category centers, then the distribution of worldly properties will appear to cluster in reliable ways. For example, focal shades of color will appear to be more common and marginal shades will appear less common, even if, objectively speaking, that is not the case. Stimuli that *in fact* vary continuously along the relevant dimension, such as a rainbow, will also *appear* to be chunked into clusters. These appearances might give subjects "empirical reason" to construct explicit representations of the relevant categories: The categories capture the way that determinate features appear to be distributed in the environment.

2.2 Vehicle-Based Category Attunement

An alternative form of implicit perceptual abstraction involves categorical sensitivity not at the level of *content* (viz., in the objects, properties, or relations represented in perception), but within the neural or syntactic *vehicles* by which those contents are represented. We call this *vehicle-based category attunement*. Vehicle-based attunement does not preclude content-based attunement, and it is possible that the two commonly co-occur. However, the two are conceptually distinct, so it is useful to distinguish them.

Vehicle-based attunement occurs when the representational vehicles used to encode particular objects or determinate features exhibit structural regularities that reliably track membership in some abstract category. Downstream processes can exploit these structural regularities to rapidly decide whether the object or feature represented by a particular vehicle belongs to the category. The relevant "structural features" might be physiological properties (e.g., rates of neural firing or population-level activity patterns), or they might be syntactic or functional-level properties (e.g., the representation's canonical decomposition into primitive constituents).

Purely for illustration, imagine a system that encodes determinate color properties using binary strings of 1s and 0s. Suppose that the strings used to encode shades within different color categories exhibit reliably different regularities. For instance, all the strings used to encode shades of red are mirror-symmetric (e.g., "110011," "101101"), all the strings used to encode shades of purple involve iteration (e.g., "101010," "010101"), and so forth (Leeuwenberg & van der Helm 2013: ch. 5). Processes that take the binary strings as input could in principle exploit these regularities to extract color category labels. This system of color representation would be set up to facilitate extraction of color categories even without content-based attunement (e.g., biases toward focal shades or "stretching" of perceptible differences along the color dimension).

One influential form of vehicle-based category attunement involves "untangling" the neural population codes through which category exemplars are encoded, enabling category membership to be determined through a simple linear decision function. Here is the idea. Think of neural activation of a particular brain region as specifying a point in a high dimensional space defined by the possible activations of all the neurons within the region. Each dimension of the space is associated with a particular neuron, with points along the dimension corresponding to its possible levels of activation. So, a point in this high dimensional space exhaustively characterizes the activation profile of the region of interest, since each point will have a determinate value along each dimension, thereby specifying a determinate activation level for each neuron in the population. Call this a "neural activation space" and consider the neural activation space defined by the retinal ganglion cells that register a retinal image. Each object category can induce innumerably many retinal images, as we vary the viewing angle, distance, adaptation of the sensory apparatus, and so on. As we vary these variables, the point in neural activation space will vary smoothly, across many—perhaps all—dimensions. Because variation can be expected to be smooth, we can conceive of the patterns of ganglion-cell activation produced by all the members of a given category as a manifold in neural activation space (see DiCarlo & Cox 2007; DiCarlo et al. 2012). As DiCarlo and colleagues see it, visual categorization is the task of distinguishing the points in neural activation space that fall within an object category's manifold, and those that do not, and then linking those points to the appropriate label.

Distinguishing the points within the manifold from others can be easier or harder, depending on the geometrical properties of the manifolds. Intuitively, manifolds can be "tangled" or "untangled". When manifolds are tangled, defining a hyperplane that neatly sorts points in one manifold from points in a different manifold is extremely difficult. If, for example, neurons in early stages of visual processing are selectively active for vertical lines, then any object that projects a vertical line on the retina will have a corresponding manifold with the same values along the dimensions corresponding to those neurons. The manifolds are therefore tangled, and deciding which object (or category) a point corresponds to is a complex matter. On the other hand, object manifolds can also be 'untangled'. Manifolds are untangled to the extent that points in the space corresponding to a given object or category all fall on the same side of some hyperplane within the high-dimensional space (see figure 3). Paradigmatically, some categories can be distinguished by linear classifiers that define a hyperplane as a simple weighted sum of neural activation. Such weighted summation computations are thought to be neurobiologically realistic (e.g., DiCarlo et al. 2012: 417). So we can think of untangled object manifolds as more useful for classification than tangled manifolds, in that they greatly simplify the decision functions needed for classification.

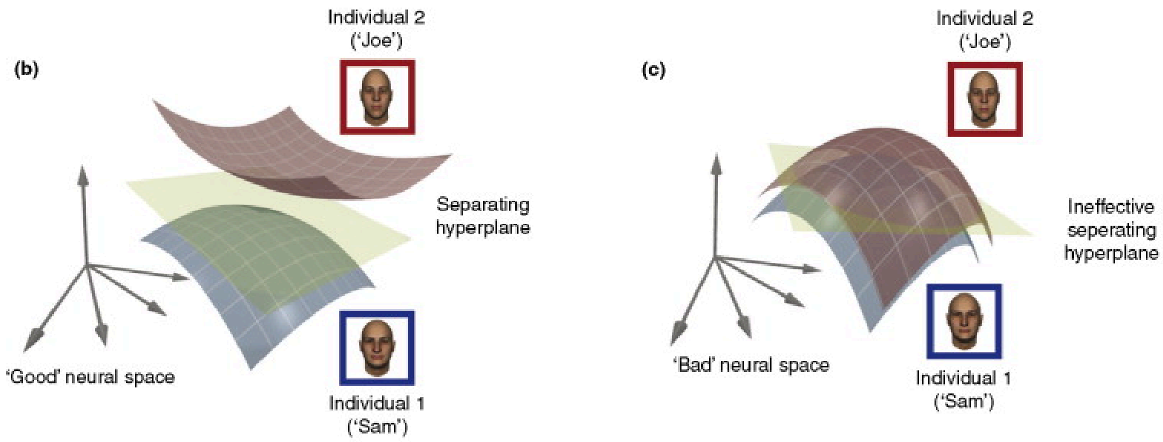

Figure 3. On the left, the patterns of activation produced by each individual fall on distinct sides of a hyperplane, making the two categories (individual 1 vs. individual 2) extractable through weighted summation. On the right, this is not the case. *Source*: DiCarlo & Cox (2007). Reprinted with permission of Elsevier.

DiCarlo and colleagues suggest that we should conceptualize the ventral visual stream as a predominantly feedforward network tasked with untangling object manifolds, beginning in V1 with highly tangled manifolds, and culminating in IT with remarkably untangled manifolds. As a result, simple linear classifiers can extract object categories from neural activation in IT, but not from neural activation in V1. Indeed, just this phenomenon has been discovered in the inferior temporal cortices of macaque monkeys (Hung et al. 2005). Though the category information is in principle recoverable from V1 activity, extracting it is extremely difficult, whereas the same information encoded in IT is relatively straightforward to extract.

Though this untangling picture may seem similar to forms of content-based attunement like perceptual magnetism, they are conceptually distinct. A location or region in perceptual similarity space comprises the *content* of a perceptual representation. Locations in neural activation space, and untangling operations defined over these spaces, have no necessary connections to representational contents. Locations in neural activation space are mathematical representations of the complete activation pattern in a region of interest. The specification is the same regardless of what—and even whether—the neural population is representing. Moreover, untangling a neural activation space *might* alter the features encoded (e.g., IT might represent category members in a canonical orientation or size, removing nuisance variation from earlier stages of vision; Buckner 2023: 113). However, the untangling process could also occur without any changes to the content represented, by simply modifying the population codes responsible for representing the very same class of contents.[7]

[7] Indeed, this is how DiCarlo and colleagues sometimes describe the process. Thus: "This conceptual framework makes clear that information is not created as signals propagate through this visual system…; rather, information is reformatted in a manner that makes information about object identity more explicit" (2012: 417).

Crucially, because the process of untangling is distinct from the process of linear classification, untangling marks a case of implicit but not explicit abstraction. The output of linear classification is a label that encodes an abstract category without encoding anything more specific-–i.e., explicit abstraction. Untangling, on the other hand, is merely a precondition for engaging in linear classification. Accordingly, perception might exhibit untangling (or other forms of vehicle-based attunement) while leaving the job of linear classification, and thus category extraction itself, to cognition. If so, then perception would display implicit but not explicit abstraction: organizing representations of determinate features in ways that grease the wheels for category extraction, without yet performing that extraction.

2.3. Does Implicit Abstraction Involve Category Representation?

Content-based and vehicle-based category attunement are ways of being differentially sensitive to abstract categories in perception. A further question concerns whether we should treat them as ways of *perceptually representing* abstract categories. In certain contexts, this question is crucial. Consider the *rich/thin* debate about perceptual content, concerning whether perception is limited to representing only "low-level" properties such as shape, size, and color, or is also capable of representing "high-level" properties such as causation, animacy, or natural kinds (e.g., Siegel 2010; Siegel & Byrne 2016; Green 2017a). If certain evidence for perceptual sensitivity to kind categories like *dog* can be explained by implicit abstraction (e.g., perceptual magnetism toward prototypical dogs), then we would like to know whether this suffices for the kind *dog* to be perceptually represented (cp. Burnston 2023). If it does, then the rich view of perception would be correct.

Perceptual representations should be posited only when they play a crucial role in psychological explanation. A key question, then, is whether we gain any explanatory power by construing cases of implicit abstraction as cases of category representation. We are skeptical that this is so. Rather, capacities for implicit abstraction seem to resemble familiar examples in which the perceptual system's architecture ensures sensitivity to some piece of information without the system needing to represent that information or draw on it during computation. To take a canonical example, the visual system might be disposed to transition from states carrying information about patterns of luminance in the retinal image to states representing the convexity of visible surfaces, and these dispositions might only make sense on the assumption that the source of light is above (Ramachandran 1988). However, absent further evidence, we see no reason to accept that "the source of light is above" is *represented* by the visual system, rather than merely being a natural constraint on its architecture (Marr 1982; Pylyshyn 2003; Orlandi 2014). Likewise, if the visual system's encoding of color is governed by a non-uniform prior peaked at category paradigms, then we might treat this as the perceptual system *respecting* the presence of privileged color categories without *representing* them.

If implicit abstraction fails to qualify as category representation, then this raises the bar for establishing a rich view of perceptual content. To show that an abstract, high-level category like *dog* is perceptually represented, it would be insufficient to show that perception is systematically sensitive to the boundary between dogs and non-dogs, or that it is sensitive to correlations among lower-level feature dimensions diagnostic of dog-hood (e.g., Burnston 2023). Rather, it would need to be shown that the category *dog* is extracted by means of a category label.

Alternatively, some might wish to treat content- and vehicle-based category attunement as cases of *implicit category representation*. Unfortunately, there is little agreement about how to understand the notion of implicit representation, with some arguing that the notion should be abandoned altogether (e.g., Ramsey 2007: ch. 5). Shea (2015) argues that a key distinction between implicit and explicit representation is that explicit representations make the information they represent available to a wide range of consuming processes while implicit representations function only within a narrow range of processes (see also Clark 1992). And indeed, many cases of implicit abstraction seem to fit the latter model. For instance, in the case of perceptual magnetism, sensitivity to color categories might be built into the processes that compute a surface's determinate shade of color. Such sensitivity is embodied in a surface's being pulled toward one perceptual magnet rather than another. However, perceptual magnetism alone would not be expected to make information about the item's color category *per se* (over and above its apparent determinate shade) available to perceptual processes outside those responsible for computing surface color. Accordingly, *if* implicit abstraction is to be understood as a kind of category representation, then it seems best construed as implicit representation.

If the reader wishes to treat implicit abstraction as a kind of implicit category representation, then we have no objection. It is also possible that on some psychosemantic theories of perceptual representation, certain cases of implicit abstraction will count as representational while others will not. For our purposes, it is enough that we recognize the distinction between implicit and explicit abstraction, where only the latter involves forming a discrete representation that encodes a category without encoding anything more specific. Whatever the proper representational treatment of implicit abstraction, it remains distinct from explicit abstraction.

Note finally that even if implicit abstraction doesn't involve genuine category representation, this does not mean that the thin view of perception is correct. After all, it is possible that perception exhibits *explicit* abstraction for certain high-level properties. We won't take a stand on the rich/thin debate here. However, we note that both sides of this debate can make use of the diagnostic criteria for explicit abstraction we propose later on (section 4) in adjudicating whether perception explicitly represents high-level properties, or is only sensitive to them in some weaker way.

2.4. Against Permissive Conceptions of Abstraction

We've argued that perception is implicitly attuned to abstract properties when perceptual representations of determinate properties are appropriately biased, structured, or organized. While this falls short of category extraction, it marks a way that perception puts us in position to think about abstract properties despite its determinacy. However, some have endorsed views on which the perceptual representation of abstract properties comes "for free" whenever we perceptually represent determinate properties. This approach makes perceptual abstraction a significantly less demanding psychological achievement, and thus a far more pervasive phenomenon. We'll argue that it should be rejected.

One way to motivate a permissive conception of perceptual abstraction is through reflection on ordinary looks-reports. It might be thought that whenever an object looks scarlet, it thereby also looks red, or even colored. More generally, whenever an object looks F, it also looks any determinable of F.[8] Someone who finds this view plausible might be attracted to a conception of perceptual representation that comports with it, according to which the perceptual representation of determinate properties suffices for the perceptual representation of their determinables.

One account delivering a result in this neighborhood is John Kulvicki's view that perception pervasively exhibits "vertically articulate" content spanning levels of abstraction (Kulvicki 2007, 2015; see also Munton 2021: 648-649). To say that perceptual representations have vertically articulate contents is to say that they represent highly specific properties as well as abstractions from those properties at various levels of determinacy.

Kulvicki uses a contrast between analog and discursive representations to explicate his view. Consider the English words for various shapes, such as "square," "triangle," "circle," etc. Though being square entails being rectangular, no features of the words "square" and "rectangle" make this information easily available to a user of the representations. By contrast, consider a mercury thermometer, which represents temperature using heights of mercury. Having a particular height represents the temperature as, say, 75 degrees fahrenheit. But in this case, the representational vehicle *also* has the property of having a height that falls in a particular range—namely, the property of having a height somewhere in the interval between the points marked "70" and "80." This more abstract vehicular property of the thermometer carries the information that the temperature is between 70 and 80 degrees fahrenheit (Kulvicki 2007, 362). Kulvicki suggests that a hallmark of analog systems is that they carry information about myriad abstractions from their most determinate contents in *extractable form*. For a representation to carry information in extractable form, there must be a property of the vehicle that carries the information, without that property carrying any more specific information. So, the property of having a mercury height between the points marked "70" and "80" is a property of the vehicle (thermometer), which carries the information that the temperature is between 70 and 80 degrees, and *that property* of the vehicle carries no more specific information. On the other hand, *no* property of the word "square" that carries the information [rectangle] without carrying any more specific information (e.g. [square]).

On Kulvicki's view, then, perceptual content is vertically articulate, and this is explained by the analog format of perceptual representation. Analog perceptual representations *represent* the wide range of abstract properties that they carry information about in extractable form. The view that all

[8] Compare Helton's (2018: 250) 'perceptual closure principle'.

perceptual representation is analog is controversial (Green 2023a). Fortunately, we can sidestep this issue for present purposes. So long as *many* dimensions are perceptually represented in analog form (cp. Beck 2019), it would follow that perception represents vertically articulate contents along the relevant dimensions.

We emphasize that *carrying information in extractable form* is distinct from *extraction*. To extract a property is to form a representation (or separable constituent of a representation) that encodes the property without encoding anything else, or anything more specific. Even if analog representations carry information in extractable form about properties more abstract than their most determinate contents, typically they don't extract those properties. A given thermometer reading of 75 degrees has a determinable vehicular property that carries the information that the temperature is between 70 and 80 degrees, but it does not extract this information because that very same representation also encodes a far more specific temperature property.

Thus, Kulvicki's view implies that perceptual abstraction comes for free whenever we have an analog system of perceptual representation. So, assuming that analog perceptual representations are rife, so is perceptual abstraction. However, we now argue that this account of perceptual abstraction is too permissive, since it fails to secure any distinctive explanatory role for the perceptual representation of abstract properties.

Note that when the thermometer registers 75 degrees, the mercury's height falls not just between "70" and "80", but within myriad other intervals (e.g., between "61" and "93," between "2" and "76," etc.). These determinable properties of the vehicle carry information about myriad abstract temperature properties. So, it follows from Kulvicki's view that the thermometer reading also represents the temperature as between 61 and 93 degrees, between 2 and 76 degrees, and so on. On a *very* thin conception of representation, perhaps this is unobjectionable. But we contend that such a conception does not help us understand how perception is selectively sensitive to certain abstract properties rather than others.

Suppose that perceptual color representations are analog insofar as they exhibit values of a vehicular magnitude M that reliably covaries with, or *mirrors*, color values. If so, then nearby values of M function to encode similar colors. Following Kulvicki, we can say that a perceptual representation's exhibiting a value of M within some interval I carries the information that a corresponding determinable color category is present: namely, the category including all and only those shades encoded by values of M within I. For any determinate value of M, there are countless such intervals to which it belongs, and only a small number of those will carry information about regions of color space associated with a single focal shade. More concretely, an analog perceptual representation of some determinate shade of red (say, burgundy) would, on Kulvicki's view, exhibit an abstract vehicular property carrying the information that the perceived object is red, but would *also* exhibit an abstract vehicular property carrying the information that the perceived object is (say) red-or-purple, or that it belongs to an alien category spanning half of the red region and half of the purple region.

Thus, if having an abstract vehicular property that carries information about a given abstract category is sufficient for that abstract category to be perceptually represented, then it follows that analog perceptual representations of burgundy simultaneously represent *all* of these abstract color categories, not just *red*.

However, the mind is wholly insensitive to most of these vertically articulate contents. Consider the prototypical categorical perception phenomenon: inter-stimulus differences crossing category boundaries are (on average) more noticeable or salient than within-category differences. In the case of color, differences that cross the red-purple boundary are more salient than objectively equivalent differences within either category (Skelton et al. 2017). As we've seen, Kulvicki's view does predict that perception represents the categories *red* and *purple*. However, on Kulvicki's view, there is nothing special about these categories. The same analog vehicle will also encode arbitrary categories that crosscut our intuitive color categories, and for which no categorical effect is found. Accordingly, the presence of vertically articulate content cannot explain these forms of categorical sensitivity in perception. Perception prioritizes certain abstract categories over others, even when both are abstractions from the most determinate properties we perceive. It puts us in position to cognize the former categories, and discourages us from cognizing the latter. An account of perceptual abstraction should explain what is special about those abstract categories to which perception is differentially sensitive, and vertically articulate content is ill-suited to do so.

Likewise, reflection on ordinary looks-reports should not be taken to motivate a permissive conception of perceptual abstraction. Suppose it's true that whenever an object looks F, it also looks any determinable of F. Then the appropriate lesson to draw, we suggest, is that "looking" to have some abstract property is insufficient for perceptually abstracting that property. We are interested in cases where perception exhibits *differential sensitivity* to abstract properties, and the perceptual system of a subject to whom a given surface looks burgundy need not exhibit any differential sensitivity to redness *simpliciter*. For example, it need not prioritize redness over the property of being red-or-purple.

Despite disagreeing with certain aspects of Kulvicki's view, we are open to the idea that analog formats are especially *amenable* to perceptual abstraction given further processing or modification. Perhaps representations in analog form are well-suited to operations that issue in either implicit or explicit abstraction. Nonetheless, it is important to distinguish the forms of abstraction that a representational format *permits* from the forms of abstraction that it *actually exhibits*. Analog format might permit indefinitely many abstractions from determinate perceptual content with appropriate further factors in place (e.g., perceptual magnetism, untangling operations, or full-fledged extraction) without exhibiting any particular attunement to those abstractions versus others.

Thus, we contend that a theoretically useful account of perceptual abstraction should require the operations of the perceptual system to be differentially sensitive to abstract properties, since only then do the relevant properties play a distinctive explanatory role in perceptual processing. Mere possession of vehicular properties that carry information about abstract properties is insufficient for perceptual abstraction in the sense we care about.

3. Why the Distinction Between Implicit and Explicit Abstraction Matters

This section considers two areas where the distinction between implicit and explicit abstraction matters: debates about the role of concepts in perception, and rationalist–empiricist debates about our innate representational repertoire. First, however, we set up the discussion with a case study in which evidence that has been taken to support explicit perceptual abstraction can be equally explained by implicit abstraction.

3.1. Categorical Perception of Color

Our case study involves perceptual sensitivity to color categories. Various studies have been taken to suggest that perceptual color differences are accentuated at category boundaries, even in early infancy. For example, Franklin, Pilling, and Davies (2005) examined the categorical perception of color in both 4-month-old infants and adults. Participants saw a colored patch against a differently colored background of the same luminance. The patch and the background either belonged to the same color category or to different color categories, but in both cases their colors were equidistant in CIE color space. Franklin et al. found that both 4-month-olds and adults were faster to fixate the colored patch when the patch and the background belonged to different color categories than when they belonged to the same category, suggesting that the patch stood out more from its background in the different-category condition.[9]

Some have sought to explain this perceptual sensitivity to color categories in terms of explicit category representations. On this view, the reason why surfaces belonging to different color categories are more perceptually dissimilar than surfaces belonging to the same color category is that the former surfaces are explicitly represented as sharing a property that the latter surfaces are not: their color category. Thus, Block (2023) writes:

> What is a perceptual category representation? One definition often given is that, for perceptual categories, discrimination across boundaries is more fine-grained (i.e., more sensitive to objective differences) than discrimination within boundaries. (…) Such definitions, though, may be less fundamental than another approach based on perceptual attributions of the categories themselves. Better discrimination across than between borders may be just an index of those categorical representations. (Block 2023: 271)

[9] For further evidence supporting sensitivity to color categories in infancy, see Bornstein et al. (1976), Franklin & Davies (2004), and Skelton et al. (2017). Adult perception displays sensitivity to color categories in visual search (Daoutis et al. 2006; Wakui et al. 2025; although see Brown et al. 2011) and multiple-object tracking (Sun et al. 2020, 2023).

Here, Block conjectures that superior discrimination across category boundaries than within categories is an *index* of *something else*—namely, category representation.[10] This is to treat category representation as something over and above the patterns of perceptual similarity themselves. Accordingly, by "categorical representation" Block cannot mean merely *implicit* abstraction. For, in implicit abstraction, category representation is not something *independent* of the relevant category-based similarity patterns. Indeed, on our preferred approach, implicit abstraction doesn't involve category representation at all. But even if one understands implicit abstraction as a kind of implicit category representation, still the relevant category representations are not independent of category-based patterns of perceptual similarity, but (at least in the case of perceptual magnetism) are simply constituted by those patterns.

However, the inference from categorical sensitivity to explicit category representation is unsound. The infant data Block cites, alongside other kinds of perceptual sensitivity to color categories, are also explainable by appeal to perceptual magnetism within color perception (see Green 2024). If representations of focal shades at the centers of color categories are perceptual magnets, then perceptual representations of determinate shades will tend to be 'pulled' towards them, resulting in subjects judging shades in a focal shade's 'magnetic field' to be more similar to the focal shade, and hence less similar to shades in other 'magnetic fields' (see Bae et al. 2015; although see Dubova & Goldstone 2022 for evidence of anti-magnetism patterns under some conditions).[11] As a result, shades in the same color category will be, on average, more perceptually similar to one another than to objectively equidistant shades in a different color category, leading colored patches to stand out more against backgrounds of a different color category than against backgrounds of the same color category, as Franklin et al. (2005) report. This hypothesis can also be used to explain parallel cases in which perception treats surfaces from different color categories as more dissimilar than objectively equidistant surfaces from the same category, such as in multiple-object tracking (Sun et al. 2020) or visual search (Daoutis et al. 2006).

So, evidence for perceptual sensitivity to categories shouldn't be taken to establish explicit perceptual representations of categories. We now highlight the significance of this point, focusing on its relevance to debates about the role of concepts in perception, and between rationalist and empiricist accounts of our innate representational repertoire.

3.2. Conceptualism vs. Non-Conceptualism About Perception

Block's appeal to infant categorical perception of color occurs in the context of an extended argument against Conceptualism about perception. Conceptualists claim that at least some

[10] For similar conclusions based on the infant evidence, see Chadha (2024: 16) and Laurence and Margolis (2024: 174). See, however, Garside et al. (2025) for evidence against sensitivity to color categories in non-human primates.

[11] This evidence also challenges one of Block's alleged core distinctions between perceptual and working-memory representation. Block (2024: 5) contends that working-memory representations, but not perceptual representations, are biased toward category centers. However, the most plausible interpretation of the evidence is rather that biases toward category centers are *already* present in perception but simply *exaggerated* in working memory.

perceptual representation is conceptual. More precisely, at least some perceptual representations include tokens of representation-types that are also tokened and deployed within conceptual thought (e.g., logical inference). While some claim that *all* perception is conceptual (McDowell 1994; Brewer 1999), a more popular view in recent years holds that only *some* perceptual representation is conceptual (e.g., Peacocke 1992: 88; Mandelbaum 2018; Quilty-Dunn 2020; see Green 2024 for discussion). An attractive feature of this hybrid form of Conceptualism is that it evades many traditional arguments for Non-Conceptualism, including Evans's (1982) fineness-of-grain argument (see Green 2024: 7-8). Moreover, the claim that the representations available to perception and conceptual thought only *partially* overlap accommodates other data that have been wielded to support Non-Conceptualism, such as the observation that perception is allegedly incapable of representing logically complex contents (e.g., Geach 1967: 22-23; Block 2023: ch. 4). After all, the claim that perception has access to *some* concepts doesn't entail that it has access to *all* concepts. Even if perception lacks access to concepts of logical connectives or modal operators, it doesn't follow that perception never deploys concepts.

The claim that only some perceptual representations are conceptual naturally raises the question of *which* ones are conceptual. One account holds that perception deploys concepts when it engages in *categorization*. In this vein, Mandelbaum (2018) ties the issue of whether perception deploys concepts to the issue of whether perception exhibits categorization. He writes:

> As I will use the term, "conceptualists" are theorists who maintain that categorizing a visually presented stimulus (henceforth: categorizing) occurs perceptually, so that perception outputs an already conceptualized representation. In contrast, "nonconceptualists" propose that the outputs of perception are non-conceptual representations, which entails that categorization must occur post-perceptually. (Mandelbaum 2018: 267)

If, as Mandelbaum claims, Conceptualism *just is* the view that categorization occurs perceptually, then perceptual categorization must require the deployment of concepts. Call this view *Categorization Conceptualism*.

Block's argument threatens to undermine Categorization Conceptualism, and in this respect it is more powerful than traditional Non-Conceptualist arguments such as the fineness-of-grain argument. In particular, Block argues that in the case of infant color perception, perceptual category representation occurs without deploying color concepts. His evidence that infants fail to deploy concepts of color categories is that they often seem to ignore color when reidentifying objects over time or when forming expectations about how objects will appear after emerging from behind an occluder (see Wilcox 1999; Tremoulet et al. 2000).[12]

[12] However, the story from developmental psychology is far less straightforward than this characterization suggests. Other studies suggest that young infants do consult color information when reidentifying objects over time, even without any special priming (Kaldy & Blaser 2009; Bremner et al. 2013).

If infants exhibit perceptual categorization with respect to color without deploying color concepts, then this directly challenges Categorization Conceptualism. Perceptual categorization can occur non-conceptually.

However, "perceptual categorization" is ambiguous between implicit and explicit forms of abstraction. On one reading, "perceptual categorization" is just another term for categorical perception, i.e. the phenomenon where objects from the same category are (on average) more perceptually similar and less discriminable than objects belonging to different categories (Goldstone 1994; Goldstone & Hendrickson 2010). But on another reading, "perceptual categorization" denotes the formation of an explicit perceptual representation of a category, i.e. *explicit* abstraction. We contend that the most promising version of Categorization Conceptualism interprets "perceptual categorization" in the latter, more demanding sense. Proponents of this form of Conceptualism can comfortably accept that non-conceptual perceptual representations suffice for *implicit* abstraction, i.e. categorical *sensitivity*, while reserving an important role for concepts within perception. The evidence Block adduces does not put pressure on Categorization Conceptualism so-construed, because the categorical effects are arguably explicable in terms of implicit abstraction (e.g. perceptual magnetism). We'll explore Categorization Conceptualism in more detail in section 5.

A similar issue arises in other recent Non-Conceptualist treatments of perceptual categorization. For example, among his primary examples of non-conceptual perceptual categorization, Burge (2022) includes perception of familiar shape categories such as "impala-body, dog-body, reptile-body, human-body, human-face, human-hand-body-part," and so on (428). According to Burge, the claim that a perceiver has a perceptual attributive for some shape category, such as car-shape, "implies that the perceiver can discriminate bodies with car-shapes and has a psychologically significant grouping for bodies shaped like cars" (Burge 2022: 95), elaborating later that such grouping "can amount to no more than systematically treating perceivably different instances alike in certain contexts" (2022: 127). However, the notion of a "psychologically significant grouping" is problematically obscure. Psychologically significant groupings of items within a given shape category can be explained either in terms of implicit abstraction (e.g., perceptual magnetism or vehicle-based regularities) or explicit abstraction, and we saw earlier that certain evidence favors an implicit abstraction account of our attunement to familiar shape categories (van Geert & Wagemans 2023). If the relevant categorical sensitivity, or psychological grouping, is only implicit, then Categorization Conceptualists can grant a non-conceptual treatment of them. More generally, the dispute between Non-Conceptualists and Categorization Conceptualists can only be properly adjudicated in light of the specific representational mechanisms underlying categorical sensitivity in perception.

Nonetheless, the availability of implicit abstraction as an explanatory construct places a different sort of pressure on Categorization Conceptualism. If Categorization Conceptualists claim that concepts are required only for explicit perceptual abstraction and not for implicit abstraction, then they incur an obligation to show that explicit perceptual abstraction actually occurs, and that the capacity for explicit perceptual abstraction correlates with concept possession. However, we've seen that implicit abstraction suffices to explain impressive forms of categorical sensitivity in perception. Indeed, the standard criteria for categorical perception (better discriminability across than within categories) are insufficient to establish explicit abstraction. Furthermore, vehicle-based category attunement might enable us to rapidly cognize abstract categories (such as *dog*, *flowers*, or *beach scene*) on the basis of perception without forming explicit representations of these categories within the perceptual system. This fact makes it difficult to determine the representational underpinnings of our capacities to detect the presence of basic-level categories on the basis of very brief stimulus presentations (e.g., Potter et al. 2014). On one view, perception explicitly represents basic-level categories and transmits these representations to cognition (Fodor 1983; Mandelbaum 2018). On another view, feedforward perceptual processes function to structure or "reformat" representations of individual category members in a manner that allows basic-level categories to be extracted through rapid computational procedures such as linear classification (DiCarlo & Cox 2007), but leave category extraction itself up to post-perceptual cognition.

But if so much perceptual categorical sensitivity can be explained by implicit abstraction, then what *would* count as good evidence for explicit perceptual abstraction? Categorization Conceptualists need an answer to this question, and we return to it in section 4. First, however, we develop another application of the distinction between implicit and explicit abstraction.

3.3. Rationalist vs. Empiricist Accounts of the Acquisition Base

Our focus has been on abstraction construed as a psychological achievement rather than as a learning or acquisition process. Nonetheless, a nuanced understanding of abstraction in our sense has implications for the debate about abstraction construed as an acquisition process. On this latter usage, "abstraction" refers to the Lockean process of acquiring representations of relatively less determinate properties on the basis of representations of more determinate properties, e.g. coming to represent rectangularity on the basis of representing a number of determinate rectangles. For clarity, in this subsection we'll call the relevant acquisition process "acquisition-abstraction," and the psychological achievement "achievement-abstraction." We'll revert to using "abstraction" to mean achievement-abstraction in the next section. The debate about acquisition-abstraction has predominantly been cast as between rationalists and empiricists, and concerns the origin of abstract ideas or concepts.[13]

What is the character of the psychological traits that underpin the emergence of abstract concepts? Empiricists characteristically hold that such concepts are acquired by domain-general learning mechanisms, with a relatively restricted set of representational primitives, while rationalists characteristically posit domain-specific learning mechanisms and relatively rich representational primitives. Acquisition-abstraction is meant to take relatively specific inputs, often within perception, and output relatively abstract representations. It has thus been of particular interest to empiricist

[13] See, respectively, Laurence and Margolis (2012; 2024, Chapter 5) and Buckner (2018; 2024, Chapter 3) for contemporary rationalist and empiricist perspectives on acquisition-abstraction.

theorists, since the more representations that can be explained by acquisition-abstraction, the fewer representational primitives are required in the acquisition base.

The distinction between implicit and explicit achievement-abstraction enables a nuanced taxonomy of views along the spectrum from Empiricism to Rationalism. Specifically, it allows us to delineate ways that our innate (i.e., unlearned) representational endowment might bias us toward acquiring certain concepts without yet *containing* those concepts. We illustrate this for the acquisition of color concepts. However, the taxonomy has more general application.

At the most minimal end is what might be called 'pure empiricism', which posits no perceptual prioritization of the relevant categories in the acquisition base. This model is especially plausible for subjects learning to recognize arbitrary categories first encountered in experimental contexts.[14] For color, pure empiricism would hold that innate color representations in no way prioritize particular shades or regions of similarity space. Hence, learning color categories is relevantly similar to learning to recognize arbitrary properties based on experimenter feedback.

A second model posits a uniform similarity space, but specific determinate representations are highlighted or prioritized (cp. Laurence & Margolis 2024: 172-173). The similarity space is uniform in the sense that there is no warping or perceptual magnetism. The sense in which representations of specific properties are nonetheless 'prioritized' remains open. But, for instance, it could be that focal shades (innately) attract more attention than other shades. Attentional prioritization could facilitate acquisition-abstraction by offering especially salient exemplars that can anchor category construction. If so, focal shades will tend to 'stand out' among possible exemplars, perhaps explaining cross-cultural convergence on the 'best examples' of color categories (Regier et al. 2005; Lindsey & Brown 2021). Alternatively, focal shades might function as privileged reference points for the perceptual representation of non-focal shades. For example, burgundy might be perceptually represented in terms of its distance from unique red, but not vice versa (see also Byrne & Hilbert 2003: 15).

A third model holds that similarity space is innately warped, exhibiting implicit achievement-abstraction. Warped similarity spaces facilitate acquisition-abstraction because they make objectively uniform dimensions appear to 'cluster', rendering categories corresponding to the clusters more natural. The warping model is especially interesting because it can explain phenomena that might initially seem to implicate explicit category representation without this extra representational machinery. For example, Laurence and Margolis (2023: 173–4) take work with infants to support the presence of innate color category representations. Infants who have habituated to a determinate color look longer at a novel determinate color that spans a color category boundary than an equally dissimilar determinate color within the same category (Bornstein et al 1976; Skelton et al. 2017). However, if the similarity space is warped, then we cannot infer from equal physical dissimilarity (or even equal spacing in CIE space) to equal subjective dissimilarity. The

[14] As in Goldstone (1994), Goldstone and Steyvers (2001), or Dubova and Goldstone (2021).

appeal to implicit achievement-abstraction can explain effects such as stronger dishabituation to changes that cross color category boundaries than to changes within categories, while denying that there are innate explicit representations of color categories.

A fourth model posits innate extraction of color categories via category labels that encode nothing more specific. The claim that color *concepts* such as RED, ORANGE, or PURPLE are innate falls into this fourth category. Accordingly, even highly systematic forms of sensitivity to color categories in early infancy do not yet establish the innateness of color concepts, since many such patterns can be explained by either the second or third models, which prescind from explicit representations of color categories.

In many domains, the best explanation for concept acquisition remains an open question. The more articulated understanding of perceptual achievement-abstraction we are developing may be helpful in answering this question, both at the level of formulating hypotheses, and in clarifying when evidence distinguishes between hypotheses. Specifically, one way of distinguishing between the third and fourth models is to determine whether perceptual sensitivity to a given category is better explained by explicit abstraction or implicit abstraction. We now consider how this might be done.

4. Probing for Explicit Abstraction: The Dissimilarity Test

Implicit abstraction can explain impressive forms of categorical sensitivity within perception. As such, it is natural to ask what (if any) forms of categorical sensitivity in perception would count as good evidence for explicit perceptual abstraction. This section proposes an evidential criterion for explicit abstraction and argues that the criterion is met by the perception of certain abstract shape categories. So, while explicit perceptual abstraction is harder to establish than one might have thought, it *is* a real phenomenon.

Because categorical biases in perceptual similarity relations are insufficient to motivate explicit category representation, we advocate a stricter evidentiary standard for explicit abstraction. To mount a convincing case for explicit abstraction, we need examples where perceptual processing is differentially sensitive to a category and where such sensitivity *cannot* be explained in terms of inter-stimulus similarity relations along the dimensions relevant to category membership. In particular, it must be shown that certain perceptual processes treat same-category stimuli as relevantly alike even when they are *more perceptually dissimilar* than different-category stimuli, even after accounting for category-based perceptual warping. Call this the *Dissimilarity Test* for explicit perceptual abstraction.

For illustration, suppose that surface A has a shade of red near the red-purple boundary, surface B has a shade of red near the red-orange boundary, and surface C has a shade of purple near the red-purple boundary. Suppose further that even after the warping effects of perceptual magnetism, A is perceived as more similar to C than to B. If certain perceptual processes treat A and B as relevantly alike (e.g., by responding invariantly to both), while treating A and C as relevantly distinct (e.g., by responding differently to each), then they cannot do so merely by exploiting the relative perceptual similarities among the stimuli. Rather, such processes would exhibit sensitivity to sameness of color category *at the expense of* similarities in perceived determinate color. This is the sort of ability that we take to provide strong evidence for explicit perceptual abstraction. It indicates that category membership per se is accessed and exploited in perceptual computation, rather than remaining implicit in perceptual similarity relations or in structural regularities among the vehicles that represent determinate properties.

We view the evidential import of the Dissimilarity Test as asymmetric. *Passing* the test provides strong evidence for explicit perceptual abstraction, but *failing* the test doesn't provide particularly strong evidence *against* explicit perceptual abstraction.[15] Nonetheless, if the test is failed, then it is incumbent on the proponent of explicit abstraction to provide some alternative motivation for explicit category representation.

We are unaware of persuasive evidence that the Dissimilarity Test is passed in the case of color perception. However, we argue that the test is passed in other domains. We discuss two examples: the perception of visual word-form categories and the perception of topological categories. In both cases, perceptual sensitivity to abstract categories is arguably best explained by perceptual representations that extract those categories.

Visual word-form categories are recovered on the basis of shape features, but shared category membership often runs counter to perceptual similarities in determinate shape. For example, despite their overall greater perceptual dissimilarity, processes that explicitly categorize by visual word-form should treat "beach" and "BEACH" as relevantly alike, and "6each" as relevantly distinct, for example by responding the same way to the former two stimuli and differently to the third. Interestingly, there is both physiological and behavioral evidence for perceptual sensitivity to shared word-form category across such variation. First, neuroimaging evidence indicates that patterns of activity in the visual word-form area (VWFA), a region of the left fusiform gyrus, are selective for individual word-forms and invariant across changes in case or font (Dehaene et al. 2001, 2004; Vin et al. 2024). For example, if "RAGE" is briefly presented and masked, then VWFA responses to "rage" are selectively reduced—a phenomenon known as neural adaptation (Dehaene et al. 2004). There is also psychophysical evidence that these patterns of neural selectivity have perceptual consequences. For example, Hanif et al. (2013) first adapted participants to an unambiguous word form, such as "cover," for 5 seconds followed by a mask. Next, they were shown a probe stimulus halfway between the previous word-form category and a distinct word-form (e.g., a stimulus halfway between "cover" and "voice") and asked them to choose which word the probe resembled more. They found that subjects' responses were biased away from the word-form category of the adapting

[15] For example, while we have argued that standard patterns of categorical perception can be explained without appeal to explicit perceptual abstraction, it is possible that categorical perception sometimes results from explicit perceptual abstraction (cp. Bae et al. 2015). The most plausible explanation of a given pattern of warped perceptual similarity might appeal to explicit abstraction, even if the relevant perceptual capacities do not pass the Dissimilarity Test.

stimulus, even if the adapting stimulus was in lowercase while the probe was constructed from its uppercase counterpart (or vice versa). So, there is some evidence that word-form categories exert adaptation aftereffects, and that these aftereffects survive changes in case (likewise, the effects appear to survive changes in handwriting style).[16] Finally, there is evidence that object-based visual priming of word-forms survives changes in case, and cannot be explained in terms of general semantic priming (Gordon & Irwin 1996; cp. Green & Quilty-Dunn 2021).

It is unlikely that visual sensitivity to word-forms can be explained by perceptual magnetism within the space of determinate shapes, since shared word-form category often comes at the expense of perceptual shape similarity. However, while we find it plausible that word-form categories are explicitly represented in the visual system, the above evidence doesn't reveal the extent to which word-form categories are actively used in perceptual processing. The case for explicit perceptual abstraction is stronger if it can be shown that explicit category representations are not just present at some early stage of processing, but also accessed by canonical perceptual computations. Thus, we turn to a second example.

Our next example involves the perception of *topological* categories. Shape properties are classified as more or less abstract based on the range of transformations under which they remain invariant, with more abstract properties preserved under a wider range of transformations than less abstract properties (Brannan et al. 2012; Todd & Petrov 2022). Within this scheme, *metric* or *Euclidean* properties are the least abstract, and are preserved only under rigid transformations and uniform scaling. They include the property of being a perfect circle or perfect square. *Affine* properties are preserved under all linear transformations and translations—for instance, stretching or shearing a figure along an arbitrary direction. Topological properties are the most abstract, and are preserved under all continuous transformations, including stretching, bending, or twisting. They include the property of being a closed figure, or of having exactly one hole (or exactly two holes, etc.). From the standpoint of topology, a coffee mug is equivalent to a donut, since both belong to the category of closed, one-holed figures.

A productive strand of research in vision science investigates our perceptual sensitivity to abstract shape properties (see Chen 2005; Bennett 2012; Todd & Petrov 2022; Green 2017b, 2023b). Studies have shown that changes that disrupt an abstract shape property are more salient than equal-magnitude changes that only disrupt metric features. Thus, a given shape change is more noticeable if it takes a figure with parallel sides to a figure with non-parallel sides than when it involves a change of the same magnitude between two figures with non-parallel sides; likewise for adding (or deleting) a corner, hole, or intersection to a shape (e.g., Chen 1982; Todd et al. 1998, 2014; Kayaert & Wagemans 2010; Amir et al. 2012; Kanbe 2013; Lovett & Franconeri 2017; Yousif & Brannon 2024). Most importantly for our purposes, there is evidence that certain perceptual processes prioritize sameness of topological category even when this patently runs counter to perceived similarities in determinate shape.

[16] See Gross (2023) for further discussion of these data.

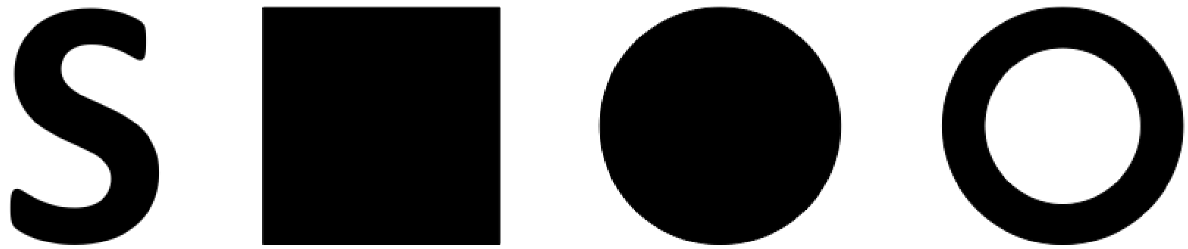

Figure 4. The left three shapes belong to the same topological category, while the rightmost shape differs in topological category.

In figure 4, the solid disc and the ring are topologically distinct, while the disc, the square, and the S-shape are topologically equivalent. The latter three belong to the same category of closed, solid figures. Despite this, the disc and ring are more perceptually similar in determinate shape than the disc and the square or the disc and the S. Furthermore, the disc is at least as different from the S as it is from the ring along a variety of other low-level measures (e.g., area, spatial frequency, and luminous flux). However, the remarkable fact is that certain canonical perceptual processes prefer to group the disc with the square and the S rather than with the ring.[17] We discuss three examples.

Apparent motion. Chen (1985) created an apparent motion display where a central figure appeared for 150 ms, vanished, and then two new figures appeared on either side after a 20-30 ms interval. One figure preserved the topological category of the original, while the other did not. Participants were tasked with reporting the direction of apparent motion, and were strongly biased (over 80%) toward perceiving motion in accordance with shared topology, even when this required large, counterintuitive changes in determinate shape, such as seeing a solid disc to morph into a solid square over a hollow ring, or a hollow square morphing into a hollow ring rather than a solid square. See also Bedford and Mansson (2010) for complementary findings.

Multiple-object tracking. Zhou et al. (2010) gave participants a multiple-object tracking task in which they tracked four moving objects amidst four distractors and then clicked on the targets at the end of the trial. The targets and distractors changed features unpredictably during tracking, and the changes either disrupted topology (e.g., from a solid disc to a ring) or left topology intact (e.g., from a disc to a square, disc to an S-shape, or disc to a dumbbell). Crucially, Zhou et al. found that only topological changes induced significant impairments to tracking accuracy. Thus, the processes responsible for computing object continuity in multiple-object tracking were selectively sensitive to changes in topological category and not to non-topological changes, even when the latter involved patently greater perceptual differences in determinate shape.

[17] Note that all four of these figures satisfy the criteria for being a "body" or "Spelke-object" (Spelke 1990; Carey 2009). While the perception of topology may be involved in perceiving Spelke-objects, the perceptual representation of topological category is more fine-grained than the perceptual representation of Spelke-objects.

Perceptual crowding. Perceptual crowding occurs when it is more difficult to recognize an object or identify its features when the object is closely flanked by other items. Crowding effects are exacerbated when the target perceptually groups with the flanking distractors based on featural similarities (Kennedy & Whitaker 2010). Xi et al. (2020) investigated whether crowding effects in peripheral vision depend on topology. Participants reported the orientation of a target flanked by distractors from either the same (e.g., S with solid disc and square) or different (e.g., S with hollow square and ring) topological categories. Crowding effects were ameliorated when the target and flankers differed topologically, suggesting that perceptual crowding mechanisms favor grouping topologically equivalent over topologically distinct objects, even when the former are perceptually more dissimilar in other respects.[18]

We contend that the perception of topological categories passes the Dissimilarity Test for explicit abstraction. Canonical perceptual processes treat items as relevantly alike when they share topological category (e.g., by grouping them together or responding invariantly to them) and as relevantly unalike when they differ in topological category, even when the former items are more perceptually dissimilar in determinate shape. To do so, the perceptual system must form representations that prioritize the respects in which a solid disc is like a solid triangle but distinct from a ring. Discrete representations of topological categories achieve this, while representations of determinate shape properties do not, since the latter leave topological categories implicit alongside information about the various respects in which a disc is more like a ring than a solid triangle (e.g., outline curvature). Moreover, the fact that topological representations are seamlessly integrated into perceptual processing strongly suggests that they are perceptual representations, rather than creatures of post-perceptual judgment. And finally, the fact that topological categories are actively used in perceptual processing suggests that they are not merely left implicit in structural regularities over perceptual vehicles (as in vehicle-based category attunement), but are actively extracted by the perceptual system and brought to bear on the relevant computations.

It is likely that the visual system needs to perform further computations beyond the formation of determinate shape representations in order to make topological categories available to perceptual processing. Because of their inherently global nature, topological properties provably cannot be computed by any set of local procedures that each depend only on a fixed number of points (Minsky & Papert 1969: 12-14). The rival view that topological categories are somehow implicit in perceptual representations of determinate shape, and that this suffices to make them accessible in perceptual processing, is computationally unrealistic.

Thus, perceptual sensitivity to topological category is a phenomenon over and above perceptual sensitivity to similarities in determinate shape. An attractive explanation is that topological categories are extracted via distinct representations from those responsible for encoding determinate shape. These representations make explicit the respects in which topologically equivalent objects are alike,

[18] There is also an advantage for detecting topological over non-topological changes to objects in peripheral vision, even when peripheral items are magnified to render non-topological properties clearly perceptible (Wu et al. 2021).

leaving out the myriad respects in which they differ. And the representations allow topological equivalence to be prioritized over similarity in determinate shape in processes such as perceptual grouping, object continuity, and apparent motion perception. Thus, not only is perception *sensitive* to topology; it *actively categorizes* by topology.

5. Perceptual Abstraction and Perceptual Concepts

Section 3.2 introduced a form of Categorization Conceptualism according to which implicit perceptual abstraction is possible without the use of concepts, while explicit perceptual abstraction requires the use of concepts. However, we also flagged a challenge for this view. To be viable, it requires that perception in fact achieves explicit abstraction. Absent strong evidence that explicit perceptual abstraction occurs, Categorization Conceptualism would be left without empirical support.

Now, however, we have argued that explicit perceptual abstraction does occur in the shape domain: Perceptual sensitivity to topological categories, and likely also to word-form categories, is not explicable in terms of warped patterns of perceptual similarity among determinate shape properties.[19] Rather, it seems best explained by the perceptual extraction of categories like *closed, solid figure*, or *closed, one-holed figure*. We think that Categorization Conceptualists should maintain that these representations are conceptual—more precisely, that they are tokens of the same representation-types that are deployed when we think about word-form or topological categories.

There are several reasons to think that explicit perceptual abstraction involves concepts. First, extraction of abstract categories is a hallmark of the conceptual. The concept CAT is used to classify an item as a cat while remaining silent about any of its more specific features, even if these features are also mentally represented. Furthermore, the view that perceptual representations of abstract categories belong to the same type as those deployed in thought has virtues of both parsimony and efficiency. Regarding parsimony, if representations of a given category are *needed* both for perceptual and cognitive tasks, then the most parsimonious view would be that the mind simply reuses the same representational resources for both purposes. And regarding efficiency, as others have noted, the claim that perceptual representations consist partly of representations of the same type as those employed in thought helps explain how the deliverances of perception can be rapidly used within cognitive processing without needing to convert them into a distinct representational format (e.g., Fodor 1983: 40; Mandelbaum 2018; Quilty-Dunn 2020; Hafri et al. 2023).[20]

[19] Of course, explicit abstraction might also occur for other categories, such as natural kinds, causation, or abstract relations (e.g., Siegel 2010; Green 2017a; Mandelbaum 2018; Quilty-Dunn 2020; Hafri & Firestone 2021; Westfall 2023).
[20] One might reply that the view that perception is partly conceptual doesn't alleviate the need for a cumbersome translation between representational formats; instead, it just locates that translation earlier (within perception) rather than later (within cognition). Reply to reply: A key advantage to placing the translation within perception is that this ensures the conversion is accomplished by modular or quasi-modular (e.g., dimension-restricted) processes (Fodor 1983; Green 2020); accordingly, it can occur quickly, automatically, and reliably (cp. Brooke-Wilson 2023).

Furthermore, the category labels unveiled by the Dissimilarity Test are well-suited to the generality and systematicity characteristic of conceptual thought. Once a system has a discrete category label, there is no obvious obstacle to recombining that label with a wide variety of other representations, including logical connectives. Note that the same does not hold for implicit abstraction. Because implicit abstraction is realized by a warping of the perceptual space, or by structural regularities over myriad representational vehicles, it's not clear how the system could *token* it at all, and *a fortiori* not in conceptual thought. At minimum, tokening an entire warped similarity space—supposing for the moment that this is psychologically feasible—would not achieve the type of compact encoding of categories characteristic of conceptual thought (Quilty-Dunn 2021). After all, deploying the concept RED does not require one to call to mind every perceptible shade of red along with the perceptual similarities among them.

Finally, assuming that explicit abstraction occurs, the Non-Conceptualist owes us an explanation of why the category labels produced in perception are *not* of a syntactic type suited to participate in conceptual thought. We find the most obvious candidate explanations to be unpromising. For example, one could appeal to a strong version of modularity, on which most of the representations deployed in perception are inaccessible to central cognition. This is unpersuasive for two reasons. First, even the most ardent modularists accept that the *outputs* of perceptual systems are accessible to consuming systems in central cognition. Category labels are plausible candidates for the outputs of perception because they correspond in content to basic perceptual judgments. Second, the relevant question is whether there are token representations in perception and cognition that share a syntactic *type*. The answer to that question could be "yes" even if no token of the type is ever passed from perception to cognition. One cannot simply infer that the representations differ in syntactic type due to being deployed in different systems, both because different systems can deploy the same representation-types, and because it directly begs the question against the Categorization Conceptualist.

Alternatively, the Non-Conceptualist might argue that if the category labels produced by explicit perceptual abstraction are to function within perceptual processing, then they must be capable of integrating and composing with non-conceptual representations like maps or icons, which (allegedly) are non-conceptual on account of their lack of logical complexity or predication (Heck 2007/2023; Rescorla 2009; Camp 2018). If so, we should conclude that perceptual category labels are non-conceptual as well. However, this argument relies on the assumption that a single representation-type cannot function to compose with representations in diverse formats. But familiar examples of mixed representational formats (e.g., symbolically-adorned maps, graphs, or diagrams) and "heterogeneous inferences" integrating linguistic and map-like formats (e.g., Aguilera & Castellano 2020; Aguilera 2021) suggest that there is no such general constraint on representational composition. The mere fact that category labels *can* combine with non-conceptual representations shouldn't lead us to conclude that they can *only* combine with non-conceptual representations.

So, we think that Categorization Conceptualists should say that the representations responsible for explicit perceptual abstraction are conceptual, and there are theoretical motivations for this view. But is it *true*? While we can't definitively settle the issue here, we delve more deeply into the question in the case of topological categories. The discussion will serve to highlight the kinds of evidence that might either confirm or refute Categorization Conceptualism, and will also reveal some of the moves that Conceptualists can make in response to apparent challenges.

We first consider some *prima facie* support for Conceptualism about the perception of topological categories. Recall that Block attempts to refute Conceptualism by arguing that young infants perceptually represent color categories but fail to use them when forming expectations about objects. We argued that this argument fails because (*inter alia*) the evidence for categorical perception of color may be equally explained by implicit abstraction, and Categorization Conceptualists can simply grant that implicit abstraction may occur without concepts. But what about in the case of topology? Here, it might be wondered whether infants exhibit capacities for topological perception while failing to use information about topology in higher-level cognition. In fact, however, infants *do* seem to consult topology when reidentifying objects and framing predictions about how they will appear after emerging from behind occluders. So, when we shift focus to an area where explicit perceptual abstraction is empirically supported, we also find evidence that representations of the relevant abstract properties are available for post-perceptual processes of working memory and expectation formation.

Like adults, young infants are perceptually sensitive to sameness of topology through radical differences in determinate shape. Chien et al. (2012) found that infants aged 1.5 to 6 months who were familiarized with a solid disc subsequently dishabituated more strongly to a ring than to a solid triangle. More recently, Lin et al. (2016) extended these findings to infants aged just 0 to 4 *days*. So, the evidence suggests that perceptual prioritization of topology is present very early, and likely at birth. But when topological properties are represented in infancy, are they also available for forming expectations about objects? A study by Kibbe and Leslie (2016) suggests the answer is yes. 6-month-old infants saw a solid disc and a ring placed behind separate screens. Then, after about 2 seconds, one of the screens was raised to reveal either the same object or the other object. They found that infants looked longer when the most recently hidden object swapped topology from a disc to a ring, or vice versa. Other evidence suggests that children 3-8 years of age rely on topological properties in determining the extension of novel nouns, and often elect to extend terms based on shared topology over similarities in metric shape or color (Kenderla et al. 2023). Thus, in the case of topology, infants and young children display no sign of disconnect between the representations available to perception and those available to cognition.

More generally, Categorization Conceptualism predicts that when perception exhibits explicit abstraction for a category (which may be revealed through the Dissimilarity Test), perceptual representation of that category should correlate with the ability to use representations of the category in canonically conceptual processes such as logical inference, assuming these processes are available (we return to this qualification below). But no such prediction is made for categories that perception targets only through implicit abstraction.

We now consider *prima facie* evidence *against* the hypothesis that the perceptual representation of topological categories is conceptual. Specifically, there is evidence that perception of topological categories occurs not just in humans, but in creatures with meager cognitive capacities including honeybees and pigeons (see Watanabe et al. 2019). For example, Chen et al. (2003) trained honeybees by rewarding them for flying to one shape over another, and found that the bees were quicker to learn the rewarded shape when it differed topologically from the unrewarded shape (e.g., S-shape versus ring) than when the two were topologically equivalent (e.g., solid disc versus ring). Moreover, the bees generalized their preferences to new stimulus pairs on the basis of topological equivalence. For example, if a bee was reliably rewarded for visiting a ring over an S-shape, then it immediately preferred to visit a hollow diamond over a cross-shape.

The challenge is this: (1) If the perceptual representation of topological categories is conceptual, then such representation should only occur in creatures who have concepts of topological categories. (2) Honeybees perceptually represent topological categories while lacking concepts of those categories. Thus: (3) Perceptual representation of topological categories is not conceptual.

Categorization Conceptualists might take issue with either (1) or (2). Starting with (2), the question of whether honeybees possess concepts is likely to hinge on contentious issues about the conditions for concept possession—for instance, whether capacities for compositional, systematic, and productive thought suffice for concept possession, or whether a capacity for logically structured thought is also required (e.g., Beck 2012). As Carruthers (2009) observes, honeybees may well meet the former, less demanding conditions for concept possession. Thus, (2) is far from unassailable. However, for present purposes we want to emphasize a different line of reply available to the Categorization Conceptualist, which targets (1). The Conceptualist might maintain that explicit perceptual category representations are conceptual in *us* (or, more generally, in creatures who have capacities for logically structured thought) but *not* in honeybees.[21] Suppose that concept possession constitutively involves possessing capacities for logical inference (Beck 2012; Block 2023). If so, then creatures without the cognitive machinery for logical inference (*arguendo* honeybees) cannot possess concepts. But, on the supposed view, whether a representation-type is a concept depends on extrinsic features of the cognitive system in which it is embedded. If so, then distinct tokens of the same representation-type could differ in whether they are concepts, depending on which cognitive system embeds them. More concretely, perhaps when perception engages in explicit abstraction, the discrete category labels it produces are *of the proper form* to participate in canonical conceptual processes like logical inference. Whether they *in fact* participate in these processes depends on the cognitive environment in which they are deployed, and specifically on whether the creature also has concepts of logical connectives and appropriate combinatorial capacities. Because human psychology provides the right sort of cognitive environment for representations of topological

[21] See also Quilty-Dunn (2020: 283).

categories to play this functional role, the representations rise to the level of concepts. In creatures for whom this is not the case, they do not.

Categorization Conceptualists are therefore committed to the claim that the honeybee's representations of topology (*inter alia*) *would* be able to participate in processes of logical inference were they present in a human-like cognitive environment. It might be objected that such counterfactuals are impossible to evaluate.[22] However, we think that matters are not so bleak. If representations of topological properties play the same functional role within perceptual processing in both humans and honeybees (e.g., functioning in the same way within motion perception, tracking, Gestalt grouping, etc.), then we would have reason to posit a common representation-type in both cases. This is just an ordinary inference from sameness of functional profile to sameness of underlying representation-types. If the evidence were to reveal that this same representation-type also functions within canonically conceptual forms of thought within humans, then it would be reasonable to conclude that honeybees harbor a representation-type that, when deployed in humans, is capable of functioning in conceptual thought. Nonetheless, while we have flagged evidence that topological properties are available for *certain* post-perceptual processes (e.g., working memory and expectation-formation), it remains to be seen whether they also participate in logically structured thought. These are matters for future investigation.

Summing up: The distinction between implicit and explicit perceptual abstraction allows us to articulate a version of Conceptualism about perception that is both theoretically motivated and on firmer empirical ground than certain alternative forms of Conceptualism. Furthermore, the case study of topological perception suggests that explicit perceptual abstraction does occur, and that perceptual representations of topological categories in infancy coincide with the ability to use those categories in framing expectations about objects. Finally, while the presence of explicit perceptual abstraction in cognitively primitive creatures may seem to put pressure on Categorization Conceptualism, we have sketched ways the Conceptualist might try to accommodate this evidence.

More generally, improved clarity about the distinction between implicit and explicit abstraction advances the Conceptualism debate in fruitful ways. Armed with the Dissimilarity Test, we can better identify the battlefields on which the debate should be fought. Conceptualists should seek evidence for explicit abstraction, and Non-Conceptualists should aim either to undercut the case for explicit abstraction, or to identify cases of conceptually-endowed creatures who display explicit perceptual abstraction for a category while lacking concepts of that category. We contend that such cases have yet to be identified, and thus that Categorization Conceptualism is a viable option on the evidence.

[22] See Block (2023: 172-173).

6. Conclusion

Perception is attuned to abstract properties. Across various sensory domains, such as shape, color, and phonemes, individuals from the same category are both more perceptually similar and less discriminable than individuals from distinct categories. Sameness of category membership also influences paradigmatic perceptual capacities, such as the perception of motion, object continuity, and Gestalt grouping. Such category attunement is unsurprising, since a core function of perception is to enable us to form judgments about the categories of objects, such as *car*, *table*, or *red*. Nonetheless, perception is also *determinate*; we never perceive an object as simply *quadrilateral*, but always as some more specific type of quadrilateral. So how, despite its determinacy, does perception enable us to immediately cognize abstract categories?

We have explored multiple ways in which perception exhibits differential sensitivity to abstract properties. We have developed a distinction between *implicit* and *explicit* perceptual abstraction, distinguishing two subvarieties of the former. Implicit abstraction occurs when perceptual representations of determinate properties are biased, structured, or organized in ways that produce sensitivity to category boundaries, but without explicitly representing the categories themselves. Explicit abstraction occurs when perception produces a discrete category label—a separate representation of an abstract property over and above representations of its less abstract counterparts. Both are ways that perception puts us in a position to think about abstract categories. However, the distinction is sometimes obscured, with evidence that only supports implicit abstraction being taken to establish explicit abstraction. Moreover, the distinction is also important for longstanding debates in philosophy and cognitive science.

The distinction between implicit and explicit perceptual abstraction highlights an underexplored variety of Conceptualism about perception. Theorists on both sides of the Conceptualism debate have often treated perceptual categorization as a homogeneous phenomenon, with Conceptualists arguing that capacities for perceptual categorization suffice for perception to be conceptual (e.g., Mandelbaum 2018) and Non-Conceptualists attacking this view by arguing that perceptual attunement to abstract categories can occur non-conceptually (e.g., Block 2023; Burge 2022). However, this exchange overlooks an important possibility—namely, that only explicit perceptual abstraction requires concepts, while implicit perceptual abstraction does not. Specifically, concepts enter the picture when perception explicitly chunks underlying perceptual similarity spaces via a limited range of discrete category labels. Because the evidence suggests that explicit perceptual abstraction occurs, and that it can be empirically disentangled from implicit abstraction, this form of Conceptualism is both empirically tractable and worthy of further investigation.[23]

[23] Portions of this material were presented at Cornell University, Johns Hopkins University, Institut Jean Nicod, and the University of Glasgow. We thank these audiences for valuable feedback. Thanks also to an anonymous reviewer for very helpful comments.

References

Aguilera, M. (2021). Heterogeneous inferences with maps. *Synthese*, *199*(1), 3805-3824. doi: https://doi.org/10.1007/s11229-020-02957-w.

Aguilera, M., & Castellano, F. (2020). Maps, language, and the conceptual–non-conceptual Distinction. *Grazer Philosophische Studien*, *98*(2), 287-315. doi: https://doi.org/10.1163/18756735-00000119.

Amir, O., Biederman, I., & Hayworth, K. J. (2012). Sensitivity to nonaccidental properties across various shape dimensions. *Vision Research*, *62*, 35-43. doi: https://doi.org/10.1016/j.visres.2012.03.020.

Bae, G. Y., Olkkonen, M., Allred, S. R., & Flombaum, J. I. (2015). Why some colors appear more memorable than others: A model combining categories and particulars in color working memory. *Journal of Experimental Psychology: General*, *144*(4), 744-63. doi: 10.1037/xge0000076.

Beck, J. (2012). Do animals engage in conceptual thought? *Philosophy Compass*, *7*(3), 218-229. doi: https://doi.org/10.1111/j.1747-9991.2011.00463.x.

Beck, J. (2019). Perception is analog: The argument from Weber's Law. *The Journal of Philosophy*, 116(6), 319-349. doi: https://doi.org/10.5840/jphil2019116621.

Bedford, F. L., & Mansson, B. E. (2010). Object identity, apparent motion, transformation geometry. *Current Research in Psychology*, 1(1), 35-52. doi: https://doi.org/10.3844/crpsp.2010.35.52.

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

Bornstein, M. H., Kessen, W., & Weiskopf, S. (1976). Color vision and hue categorization in young human infants. *Journal of Experimental Psychology: Human Perception and Performance*, *2*(1), 115-129. doi: 10.1037//0096-1523.2.1.115.

Brannan, D. A., Esplen, M. F., & Gray, J. J. (2011). *Geometry*. Cambridge University Press.

Bremner, J. G., Slater, A. M., Mason, U. C., Spring, J., & Johnson, S. P. (2013). Trajectory perception and object continuity: Effects of shape and color change on 4-month-olds' perception of object identity. *Developmental Psychology*, *49*(6), 1021-1026. doi: 10.1037/a0029398.

Brewer, B. (1999) *Perception and Reason*. Oxford: Oxford University Press.

Brooke-Wilson, T. (2023). How is perception tractable?. *Philosophical Review*, *132*(2), 239-292. doi: https://doi.org/10.1215/00318108-10294422.

Brown, A. M., Lindsey, D. T., & Guckes, K. M. (2011). Color names, color categories, and color-cued visual search: Sometimes, color perception is not categorical. *Journal of Vision*, *11*(12), 2-2. doi: https://doi.org/10.1167/11.12.2.

Buckner, C. (2018). Empiricism without magic: Transformational abstraction in deep convolutional neural networks. *Synthese*, *195*(12), 5339-5372. doi: https://doi.org/10.1007/s11229-018-01949-1.

Buckner, C. J. (2024). *From Deep Learning to Rational Machines: What the History of Philosophy Can Teach Us About the Future of Artificial Intelligence*. Oxford University Press.

Burge, T. (2022). *Perception: First Form of Mind*. Oxford: Oxford University Press.

Burnston, D. C. (2023). How to think about higher‐level perceptual contents. *Mind & Language*, *38*(5), 1166-1186. doi: https://doi.org/10.1111/mila.12446.

Byrne, A., & Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, *26*(1),

3-21. doi: 10.1017/S0140525X03000013.
Camp, E. (2018). Why maps are not propositional. In A. Grzankowski & M. Montague (eds.), *Non-Propositional Intentionality*, 19-45. Oxford University Press. doi: https://doi.org/10.1093/oso/9780198732570.003.0002.
Carey, S. (2009) *The Origin of Concepts*. Oxford: Oxford University Press.
Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R. Lurz (ed.), *The Philosophy of Animal Minds*, 89-107. Cambridge University Press.
Chadha, M. (2024). Knowledge of universals. *British Journal for the History of Philosophy*. doi: 10.1080/09608788.2024.2409205.
Chen, L. (1982). Topological structure in visual perception. *Science*, *218*(4573), 699-700. doi: 10.1126/science.7134969.
Chen, L. (1985). Topological structure in the perception of apparent motion. *Perception*, *14*(2), 197-208. doi: https://doi.org/10.1068/p140197.
Chen, L., Zhang, S., & Srinivasan, M. V. (2003). Global perception in small brains: Topological pattern recognition in honey bees. *Proceedings of the National Academy of Sciences*, *100*(11), 6884-6889. doi: https://doi.org/10.1073/pnas.0732090100.
Chen, L. (2005). The topological approach to perceptual organization. *Visual Cognition*, *12*(4), 553-637. doi: https://doi.org/10.1080/13506280444000256.
Chien, S. H. L., Lin, Y. L., Qian, W., Zhou, K., Lin, M. K., & Hsu, H. Y. (2012). With or without a hole: Young infants' sensitivity for topological versus geometric property. *Perception*, *41*(3), 305-318. doi: https://doi.org/10.1068/p7031.
Clark, A. (1992). Presence of a symbol. *Connection Science*, *4*(3-4), 193-205. doi: https://doi.org/10.1080/09540099208946615.
Clark, A. (1993). *Sensory Qualities*. Oxford University Press.
Connolly, K. (2019). *Perceptual Learning: The Flexibility of the Senses*. Oxford university press.
Daoutis, C. A., Pilling, M., & Davies, I. R. (2006). Categorical effects in visual search for colour. *Visual Cognition*, *14*(2), 217-240. doi: https://doi.org/10.1080/13506280500158670.
Dehaene, S., Jobert, A., Naccache, L., Ciuciu, P., Poline, J. B., Le Bihan, D., & Cohen, L. (2004). Letter binding and invariant recognition of masked words: behavioral and neuroimaging evidence. *Psychological Science*, *15*(5), 307-313. doi: 10.1111/j.0956-7976.2004.00674.x.
Dehaene, S., Naccache, L., Cohen, L., Bihan, D. L., Mangin, J. F., Poline, J. B., & Rivière, D. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neuroscience*, *4*(7), 752-758. doi: 10.1038/89551.
DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, *11*(8), 333-341. doi: https://doi.org/10.1016/j.tics.2007.06.010.
DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How does the brain solve visual object recognition?. *Neuron*, 73(3), 415-434. doi: 10.1016/j.neuron.2012.01.010.
Dretske, F. (1981) *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
Dubova, M., & Goldstone, R. L. (2021). The influences of category learning on perceptual reconstructions. *Cognitive Science*, *45*(5), e12981. doi: https://doi.org/10.1111/cogs.12981.
Dubova, M., & Goldstone, R. L. (2022). Categorical perception meets El Greco: Categories

unequally influence color perception of simultaneously present objects. *Cognition*, *223*, 105025. doi: https://doi.org/10.1016/j.cognition.2022.105025.

Evans, G. (1982) *The Varieties of Reference*. Oxford: Oxford University Press.

Feldman, N. H., Griffiths, T. L., & Morgan, J. L. (2009). The influence of categories on perception: Explaining the perceptual magnet effect as optimal statistical inference. *Psychological Review*, *116*(4), 752-82. doi: https://doi.org/10.1037/a0017196.

Fodor, J. A. (1983) *The Modularity of Mind*. Cambridge, MA: MIT Press.

Folstein, J. R., Gauthier, I., & Palmeri, T. J. (2012). How category learning affects object representations: Not all morphspaces stretch alike. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(4), 807. doi: https://doi.org/10.1037/a0025836.

Folstein, J. R., Palmeri, T. J., & Gauthier, I. (2013). Category learning increases discriminability of relevant object dimensions in visual cortex. *Cerebral Cortex*, *23*(4), 814-823. doi: 10.1093/cercor/bhs067.

Folstein, J. R., Palmeri, T. J., Van Gulick, A. E., & Gauthier, I. (2015). Category learning stretches neural representations in visual cortex. Current directions in psychological science, *24*(1), 17-23. doi: 10.1177/0963721414550707.

Franklin, A., & Davies, I. R. (2004). New evidence for infant colour categories. *British Journal of Developmental Psychology*, *22*(3), 349-77. doi: https://doi.org/10.1348/0261510041552738.

Franklin, A., Pilling, M., & Davies, I. (2005). The nature of infant color categorization: Evidence from eye movements on a target detection task. *Journal of Experimental Child Psychology*, *91*(3), 227-48. doi: https://doi.org/10.1016/j.jecp.2005.03.003.

Garside, D. J., Chang, A. L., Selwyn, H. M., & Conway, B. R. (2025). The origin of color categories. *Proceedings of the National Academy of Sciences*, *122*(1), e2400273121. doi: https://doi.org/10.1073/pnas.2400273121.

Gauker, C. (2011). *Words and Images: An Essay on the Origin of Ideas*. Oxford University Press.

Geach, P. (1967) *Mental Acts: Their Content and Their Objects*. London: Routledge & Kegan Paul.

Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*(2), 178-200. doi: https://doi.org/10.1037/0096-3445.123.2.178.

Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General*, *130*(1), 116-139. doi: 10.1037/0096-3445.130.1.116.

Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 69-78. doi: https://doi.org/10.1002/wcs.26.

Gordon, R. D., & Irwin, D. E. (1996). What's in an object file? Evidence from priming studies. *Perception & Psychophysics*, *58*(8), 1260-1277. doi: https://doi.org/10.3758/BF03207558.

Green, E. J. (2017a). Psychosemantics and the rich/thin debate. *Philosophical Perspectives*, *31*, 153-186. doi: https://doi.org/10.1111/phpe.12097.

Green, E. J. (2017b). A layered view of shape perception. *The British Journal for the Philosophy of Science*, *68*(2), 355-387. doi: https://doi.org/10.1093/bjps/axv042.

Green, E. J. (2020). The perception-cognition border: A case for architectural division. *Philosophical*

Review, *129*(3), 323-393. doi: https://doi.org/10.1215/00318108-8311221.
Green, E. J. (2023a). The perception–cognition border: Architecture or format? In B. P. McLaughlin & J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind: Second Edition*, 469-93. Oxford: Blackwell. doi: https://doi.org/10.1002/9781394259847.ch26.
Green, E. J. (2023b). A pluralist perspective on shape constancy. *The British Journal for the Philosophy of Science*. doi: 10.1086/727427.
Green, E. J. (2024). Perceptual categorization and perceptual concepts. *The Philosophical Quarterly*. doi: 10.1093/pq/pqae144.
Green, E. J. and Quilty-Dunn, J. (2021). What is an object file? *The British Journal for the Philosophy of Science*, *72*(3), 665-699. doi: https://doi.org/10.1093/bjps/axx055.
Hafri, A., & Firestone, C. (2021). The perception of relations. *Trends in Cognitive Sciences*, 25(6), 475-492. doi: 10.1016/j.tics.2021.01.006.
Hafri, A., Green, E. J., & Firestone, C. (2023). Compositionality in visual perception. *Behavioral & Brain Sciences*, *46*, e277. doi: 10.1017/S0140525X23001838.
Hanif, H. M., Perler, B. L., & Barton, J. J. (2013). The visual representations of words and style in text: An adaptation study. *Brain Research*, *1518*, 61-70. doi: 10.1016/j.brainres.2013.04.036.
Harnad, S. (1987). *Categorical Perception*. Cambridge University Press.
Heck, R. K. (2000). Nonconceptual content and the "space of reasons". *The Philosophical Review*, *109*(4), 483-523. doi: https://doi.org/10.2307/2693622.
Heck, R. K. (2007/2023). Are there different kinds of content? In B. P. McLaughlin & J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind: Second Edition*, 110-30. Oxford: Blackwell.
Helton, G. (2018). Visually perceiving the intentions of others. *The Philosophical Quarterly*, *68*(271), 243-264. doi: https://doi.org/10.1093/pq/pqx051.
Hung, C. P., Kreiman, G., Poggio, T., & DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science*, *310*(5749), 863-866. doi: 10.1126/science.1117593.
Jenkin, Z. (2023). Perceptual learning. *Philosophy Compass*, *18*(6), e12932. doi: https://doi.org/10.1111/phc3.12932.
Kaldy, Z., & Blaser, E. (2009). How to compare apples and oranges: Infants' object identification tested with equally salient shape, luminance, and color changes. *Infancy*, *14*(2), 222-43. doi: https://doi.org/10.1080/15250000802707088.
Kanbe, F. (2013). On the generality of the topological theory of visual shape perception. *Perception*, *42*(8), 849-872. doi: https://doi.org/10.1068/p7497.
Kayaert, G., & Wagemans, J. (2010). Infants and toddlers show enlarged visual sensitivity to nonaccidental compared with metric shape changes. *i-Perception*, *1*(3), 149-158. doi: 10.1068/i0397.
Kenderla, P., Kim, S. H., & Kibbe, M. M. (2023). Competition between object topology and surface features in children's extension of novel nouns. *Open Mind*, *7*, 93-110. doi: https://doi.org/10.1162/opmi_a_00073.
Kennedy, G. J., & Whitaker, D. (2010). The chromatic selectivity of visual crowding. *Journal of Vision*, *10*(6), 15-15. doi: https://doi.org/10.1167/10.6.15.
Kibbe, M. M., & Leslie, A. M. (2016). The ring that does not bind: Topological class in infants'

working memory for objects. *Cognitive Development*, *38*, 1-9. doi: https://doi.org/10.1016/j.cogdev.2015.12.001.
Kuhl, P. K. (1991). Human adults and human infants show a "perceptual magnet effect" for the prototypes of speech categories, monkeys do not. *Perception & Psychophysics*, *50*(2), 93-107. doi: https://doi.org/10.3758/BF03212211.
Kulvicki, J. (2007). Perceptual content is vertically articulate. *American Philosophical Quarterly*, *44*, 357-369. https://www.jstor.org/stable/20464386.
Kulvicki, J. (2015). Analog representation and the parts principle. *Review of Philosophy and Psychology*, *6*, 165-180. doi: 10.1007/s13164-014-0218-z.
Laurence, S., & Margolis, E. (2012). Abstraction and the origin of general ideas. *Philosopher's Imprint, 12*, 1–22. http://hdl.handle.net/2027/spo.3521354.0012.019.
Laurence, S., & Margolis, E. (2024). *The Building Blocks of Thought: A Rationalist Account of the Origins of Concepts*. Oxford University Press.
Leeuwenberg, E., & Van der Helm, P. A. (2013). *Structural Information Theory: The Simplicity of Visual Form*. Cambridge University Press.
Liberman, A. M., Harris, K. S., Hoffman, H. S., & Griffith, B. C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, *54*(5), 358-368. doi: https://doi.org/10.1037/h0044417.
Lin, Y. L., Chien, S. H. L., & Hu, S. (2016). An early sensitivity for detecting differences in visual topological property in 0-to 4-day-old human neonates. *Pediatric Dimensions*, 1, 29-33. doi: 10.15761/PD.1000107.
Lindsey, D. T., & Brown, A. M. (2021). Lexical color categories. *Annual Review of Vision Science*, 7(1), 605-631. doi: 10.1146/annurev-vision-093019-112420.
Locke, J. (1690/1975). *An Essay Concerning Human Understanding*. Oxford University Press.
Lovett, A., & Franconeri, S. L. (2017). Topological relations between objects are categorically coded. *Psychological Science*, 28(10), 1408-1418. doi: 10.1177/0956797617709814.
Mackie, J. L. (1976). *Problems from Locke*. Oxford University Press.
Mandelbaum, E. (2018). Seeing and conceptualizing: Modularity and the shallow contents of perception. *Philosophy and Phenomenological Research*, *97*(2), 267-283. doi: https://doi.org/10.1111/phpr.12368.
Marr, D. (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: W. H. Freeman.
Matthen, M. (2005) *Seeing, Doing, and Knowing*. Oxford University Press.
McDowell, J. (1994) *Mind and World*. Harvard University Press.
McMurray, B. (2022). The myth of categorical perception. *The Journal of the Acoustical Society of America*, 152(6), 3819-3842. doi: 10.1121/10.0016614.
Minsky, M. L., & Papert, S. A. (1969). *Perceptrons*. MIT Press.
Munton, J. (2021). Visual indeterminacy and the puzzle of the speckled hen. *Mind & Language*, *36*(5), 643-663. doi: https://doi.org/10.1111/mila.12296.
Orlandi, N. (2014). *The Innocent Eye: Why Vision is Not a Cognitive Process*. Oxford University Press.
Peacocke, C. (1992). *A Study of Concepts*. MIT Press.

Potter, M. C., Wyble, B., Hagmann, C. E., & McCourt, E. S. (2014). Detecting meaning in RSVP at 13 ms per picture. *Attention, Perception, & Psychophysics*, 76, 270-279. doi: 10.3758/s13414-013-0605-z.

Pylyshyn, Z. W. (2003) *Seeing and Visualizing: It's Not What You Think*. Cambridge, MA: MIT Press.

Quilty‐Dunn, J. (2020). Concepts and predication from perception to cognition. *Philosophical Issues*, *30*, 273-92. doi: https://doi.org/10.1111/phis.12185.

Quilty‐Dunn, J. (2021). Polysemy and thought: Toward a generative theory of concepts. *Mind & Language*, *36*(1), 158-185. doi: https://doi.org/10.1111/mila.12328.

Ramachandran, V. S. (1988). Perceiving shape from shading. *Scientific American*, 259(2), 76-83. doi: 10.1038/scientificamerican0888-76.

Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge University Press.

Ransom, M., & Goldstone, R. L. (2024). Bias in perceptual learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, 15(5), e1683. doi: https://doi.org/10.1002/wcs.1683.

Regier, T., Kay, P., & Cook, R. S. (2005). Focal colors are universal after all. *Proceedings of the National Academy of Sciences*, 102(23), 8386-8391. doi: https://doi.org/10.1073/pnas.0503281102.

Rescorla, M. (2009). Cognitive maps and the language of thought. *The British Journal for the Philosophy of Science*, *60*, 377–407. doi: https://doi.org/10.1093/bjps/axp012.

Rosenthal, D. (2010). How to think about mental qualities. *Philosophical Issues*, 20, 368-393. doi: https://doi.org/10.1111/j.1533-6077.2010.00190.x.

Shea, N. (2015). Distinguishing top-down from bottom-up effects. In D. Stokes, M. Matthen, & S. Biggs (eds.), *Perception and its Modalities*, 73-91. Oxford University Press. doi: https://doi.org/10.1093/acprof:oso/9780199832798.003.0004.

Siegel, S. (2010) *The Contents of Visual Experience*. Oxford: Oxford University Press.

Siegel, S., & Byrne, A. (2017). Rich or thin?. In B. Nanay (Ed.), *Current Controversies in Philosophy of Perception,* 59-80. Routledge.

Skelton, A. E., Catchpole, G., Abbott, J. T., Bosten, J. M., & Franklin, A. (2017). Biological origins of color categorization. *Proceedings of the National Academy of Sciences*, *114*(21), 5545-5550. doi: https://doi.org/10.1073/pnas.1612881114.

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*(1), 29-56. doi: https://doi.org/10.1207/s15516709cog1401_3.

Sun, M., Hu, L., Fan, L., & Zhang, X. (2020). Tracking within-category colors is easier: Color categories modulate location processing in a dynamic visual task. *Memory & Cognition*, *48*, 32-41. doi: 10.3758/s13421-019-00959-9.

Sun, M., Xin, X., Ying, H., Hu, L., & Zhang, X. (2023). Categorical encoding of moving colors during location tracking. *Perception*, *52*(3), 195-212. doi: 10.1177/03010066221147120.

Sundström, P. (2019). Are sensory concepts learned by "abstraction" from experience? *Erkenntnis*, *84*(5), 1159-1178. doi: 10.1007/s10670-018-0002-z.

Swan, K., & Myers, E. (2013). Category labels induce boundary-dependent perceptual warping in learned speech categories. *Second Language Research*, *29*(4), 391-411. doi: 10.1177/0267658313491763.

Thau, M. (2002). *Consciousness and Cognition*. Oxford University Press.

Todd, J. T., Chen, L., & Norman, J. F. (1998). On the relative salience of Euclidean, affine, and topological structure for 3-D form discrimination. *Perception*, *27*(3), 273-282. doi: https://doi.org/10.1068/p270273.

Todd, J. T., & Petrov, A. A. (2022). The many facets of shape. *Journal of Vision*, 22(1), 1-1. doi: https://doi.org/10.1167/jov.22.1.1.

Todd, J. T., Weismantel, E., & Kallie, C. S. (2014). On the relative detectability of configural properties. *Journal of Vision*, *14*(1), 18-18. doi: https://doi.org/10.1167/14.1.18.

Tremoulet, P. D., Leslie, A. M., & Hall, D. G. (2000). Infant individuation and identification of objects. *Cognitive Development*, *15*(4), 499-522. doi: https://doi.org/10.1016/S0885-2014(01)00038-7.

Van Geert, E., & Wagemans, J. (2023). What good is goodness? The effects of reference points on discrimination and categorization of shapes. *Journal of Experimental Psychology: Human Perception and Performance*, *49*(8), 1180-1201. doi: https://doi.org/10.1037/xhp0001137.

Van Gulick, A. E., & Gauthier, I. (2014). The perceptual effects of learning object categories that predict perceptual goals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(5), 1307-1320. doi: https://doi.org/10.1037/a0036822.

Vin, R., Blauch, N. M., Plaut, D. C., & Behrmann, M. (2024). Visual word processing engages a hierarchical, distributed, and bilateral cortical network. *iScience*, 27(2). doi: 10.1016/j.isci.2024.108809.

Wakui, E., Mylonas, D., Caparos, S., & Davidoff, J. (2025). Categorical color perception shown in a cross‐lingual comparison of visual search. *Color Research & Application*. DOI: 10.1002/col.22964.

Watanabe, A., Fujimoto, M., Hirai, K., & Ushitani, T. (2019). Pigeons discriminate shapes based on topological features. *Vision Research*, *158*, 120-125. doi: https://doi.org/10.1016/j.visres.2019.02.012.

Westfall, M. (2023). Perceiving agency. *Mind & Language*, *38*(3), 847-865. doi: https://doi.org/10.1111/mila.12399.

Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition*, *72*(2), 125-166. doi: 10.1016/s0010-0277(99)00035-9.

Wu, R., Wang, B., Zhuo, Y., & Chen, L. (2021). Topological dominance in peripheral vision. *Journal of Vision*, 21(10), 19-19. doi: https://doi.org/10.1167/jov.21.10.19.

Xi, H., Wu, R., Wang, B., & Chen, L. (2020). Topological difference between target and flankers alleviates crowding effect. *Journal of Vision*, 20(9), 9-9. doi: https://doi.org/10.1167/jov.20.9.9

Yousif, S. R., & Brannon, E. M. (2024). Perceiving topological relations. *Psychological Science*. doi: 10.1177/09567976241309615.

Zhou, K., Luo, H., Zhou, T., Zhuo, Y., & Chen, L. (2010). Topological change disturbs object continuity in attentive tracking. *Proceedings of the National Academy of Sciences*, 107(50), 21920-21924. doi: https://doi.org/10.1073/pnas.1010919108.