DOI: 10.1111/phis.12185

ORIGINAL ARTICLE

Philosophical Issues
A Supplement to NOÛS

Concepts and predication from perception to cognition

Jake Quilty-Dunn

University of Oxford & Washington University in St. Louis

Correspondence
Jake Quilty-Dunn, Washington University in St. Louis, St. Louis, MO, 63130.
Email: quiltydunn@gmail.com

1 | CONCEPTUALISM

One popular doctrine in 20th-century philosophy was *conceptualism* about perception. The core idea was that perceptual awareness is structured by concepts possessed by the perceiver. A primary motivation for conceptualism was epistemological: perception provides justification for belief, and this justificatory relation is only intelligible if perception, like belief, is conceptually structured (Brewer, 1999; McDowell, 1994; Sellars, 1956). We perceive *that a is F*, and thereby grasp perceptual evidence that justifies the belief *that a is F* and inferentially integrates with premises like *If a is F then a is G* to produce the belief *that a is G*.

Conceptualism is less popular today (cf. Bengson, Grube, & Korman, 2011; Mandelbaum, 2018; Mandik, 2012). The *a priori* justification for conceptualism has crashed face-first into a wall of empirical evidence. For instance, children and non-human animals possess perceptual capacities despite lacking many hallmarks of conceptual cognition (Bermudez, 1998; Burge, 2010a; Block ms). Meanwhile, in adults, mental imagery and related phenomena implicate iconic rather than conceptual/propositional formats (Carey, 2009; Fodor, 2007; Quilty-Dunn, 2019a). A growing contingent of theorists thus regard perception as a natural kind marked by its proprietary nonconceptual representations (Burge, 2014; Burnston, 2017a; Carey, 2009; Kulvicki, 2015a; Toribio, 2011; Block, ms; see also Evans (1982); Hopp (2011); Peacocke (2001) for other nonconceptualist arguments).

Though opinion has shifted strongly in favor of nonconceptualism, it may be time for the pendulum to swing back. Putting the traditional normative motivations for conceptualism aside, it makes sense even from a purely descriptive, naturalistic perspective that at least some of the vehicles of perception should be conceptual. Many cognitive operations make use of concepts; thus many cognitive responses to perception would be facilitated if some outputs of perception came prepackaged in a conceptualized format.

This point fits with modularity-based accounts of perception, and was fittingly made by Fodor in his discussion of input modules as "subsidiary systems" that must "provide the central machine with information about the world; information expressed by mental symbols in whatever format cognitive processes demand of the representations that they apply to" (Fodor, 1983, p. 40). Similarly, Mandelbaum argues that the outputs of modular perceptual systems ought to be conceptualized in order to "actually guide action by entering into other cognitive processes" (2018, p. 271). It is an underemphasized explanatory virtue of modularity that it allows for a system to be distinctly perceptual (in virtue of its modularity) while outputting representations that are immediately consumable by cognition (in virtue of their format). Modularity-based versions of conceptualism thereby avoid full-fledged versions of the "interface problem" in interactions between perception, cognition, and action (Burnston, 2017b; Butterfill & Sinagaglia, 2014; Mylopoulos & Pacherie, 2017; Shepherd, 2018; 2019).

It is fully compatible with this modularity-based conceptualism that some perceptual processes output representations in nonconceptual (e.g., iconic) formats. Instead of insisting on conceptual structure as a transcendental epistemological requirement, modularity-based conceptualists can be pluralists about perceptual representation (Quilty-Dunn, 2019b). As long as some significant component of perception is conceptual and feeds immediately into cognition, there is room for other perceptual representations to have other formats with other functional advantages. For example, perhaps iconic representations allow for richer, messier content to be encoded in perception, while sparse conceptual representations provide neatly packaged categorizations to central cognition.

However, perception is older than cognition. One might object that our perceptual systems evolved from creatures who lacked cognition, therefore there was no evolutionary pressure for concepts to figure in perception. In what follows, I'll sketch a version of conceptualism that posits concepts in perception independently of stimulus-independent cognitive abilities. In particular, I'll argue not only that adult humans have conceptually structured perceptual representations, but also that these conceptual outputs of perception constitute a natural representational kind found in children and animals alike. Perceptual object representations function to segment out particulars, track them, and predicate features of them, including conceptual categories. These object representations constitute an evolutionarily ancient and developmentally early source of predicate-argument propositional structure that is useful for (1) tracking individuals, (2) subsuming them under categories, and (3) distinguishing reference-guiding elements from pure attributions. These structures can function as evidential inputs to inferential processes in creatures that have the requisite inferential abilities.

I will first argue against stimulus-independence as a constitutive condition on conceptuality (Prinz, 2002, p. 197; Beck, 2018; Burge, 2010b; Camp, 2009) in favor of a Cartesian view that concepts are simply representations of a certain sort that, in principle, require no particular mental abilities for their instantiation in human and animal minds (Fodor, 2004). I'll then use empirical evidence to argue that, in fact, perceptual object representations are conceptualized propositional structures that develop (and likely evolved) prior to creatures' abilities to use them in inference. The resulting picture preserves much of the letter—if not exactly the spirit—of traditional conceptualism.

2 | INDIVIDUATING CONCEPTS

2.1 | Stimulus-independence

It is not entirely clear how we ought to understand stimulus-dependence. Lots of mental activity might happen to be prompted by a pattern of stimulation and happen to end when stimulation ends, but this wide net might capture a messy variety of mentation rather than a natural kind. One could reasonably demand well-defined, testable characterizations of stimulation and dependence thereupon, and difficulties will surely arise in trying to provide them (see Beck (2018) for careful discussion of the details). I'll discuss two forms of stimulus-independent use: recombinability and logical inference. However, I propose to grant in general that there is some notion of stimulus (in)dependence that's coherent enough to figure in a candidate condition on concept possession. What matters for present purposes is the following claim: *concepts are the sorts of mental phenomena that can only occur in creatures that have the ability to deploy them independently of what their transducers are doing at the moment*.

I'll also put aside a particularly strong form of the claim at issue. Beck (2018) argues that perception and cognition are distinguished by means of stimulus-independence: a state/process is perceptual iff it is stimulus-dependent, and cognitive iff stimulus-independent.

This formulation runs into a counterexample: perception-based demonstrative thought, which is stimulus-dependent but cognitive (Beck, 2018, pp. 328–329). Beck's way of responding to this counterexample is to add the condition that *every element* of a representation must be stimulus-dependent for the representation (and the process that produces it) to be perceptual (2018, p. 330). Since demonstrative thoughts have concepts as elements—e.g., the concept RED is the predicative element in the thought THAT IS RED—and concepts are redeployable elsewhere, demonstrative thoughts fail this additional criterion. However, this additional criterion simply rules out the possibility of deploying concepts in perceptual systems by fiat. It seems like a largely empirical question whether humans can deploy concepts perceptually (Jacobson & Putnam, 2016; Mandelbaum, 2018). Our theories shouldn't build in the impossibility of this scenario to avoid counterexamples.

One available move for Beck's view would be to distinguish states and processes: concepts are stimulus-independent states, but they can be deployed via stimulus-dependent perceptual processes. On this relaxed condition, however, perception-based demonstrative thoughts again represent a counterexample. Thus Beck's version of the stimulus-dependence criterion faces a dilemma: avoid demonstrative thought as a counterexample but render perceptual deployment of concepts impossible by fiat; or allow for the perceptual deployment of concepts but succumb to the counterexample.

My main target is not the thesis that perception should be analyzed in terms of stimulus-dependence, but instead the hypothesis that having concepts constitutively requires the ability to use them in a stimulus-independent way—call that hypothesis STIMULUS-INDEPENDENCE. It's compatible with STIMULUS-INDEPENDENCE that perception is constituted by something completely unrelated, such as proprietary representational formats or modularity. It's even compatible with this hypothesis (so stated) that concepts could be deployed via stimulus-dependent processes. As long as a creature with the concept RED can deploy that concept independently of stimulation, it's entirely possible that the creature might have some stimulus-dependent means of deploying it as well.

Though STIMULUS-INDEPENDENCE allows for concept deployment in perception, it places significant constraints on the kinds of creatures that can deploy concepts. In particular, such creatures must possess the ability for stimulus-independent thought. A member of a species that has evolved perceptual systems and uses them to guide action but lacks central cognition cannot have concepts. Such creatures are "passive reactors, at the mercy of their environments" (Camp, 2009, p. 290), lacking the cognitive freedom that marks conceptual thought. Thus, for Camp, the representations they deploy in perception and action-guidance must be nonconceptual. For a state "to even be a candidate for being conceptual, it must be cognitive" (2009, p. 279). Specifically, concepts are "cognitive, representational abilities that are systematically recombinable in an actively self-generated, stimulus-independent way" (Camp, 2009, p. 302).

Likewise, for Burge, a representation is conceptual iff it "can function in pure predication" (2010b, p. 45), which requires functioning "outside the scope of a context-bound identificational, referential structure" (2010b, p. 44). For example, in 'That table is brown', the predicate 'table' functions within the scope of a referential noun phrase ('That table'), while 'brown' functions purely predicatively. Burge assumes that perceptual contents exclusively have context-bound identificational, referential structures (e.g., That F). Perception therefore cannot suffice for pure predication. Since possessing concepts requires the ability to use them in pure predication, possession of concepts requires the ability to use concepts outside perceptual contexts. For Burge, a paradigm example of such use is logical inference, discussed at length below. Thus Burge's view leads naturally to a version of STIMULUS-INDEPENDENCE.

A large part of the intuitive appeal of STIMULUS-INDEPENDENCE is that it captures what Kant called the "spontaneity" of thought (Kant, 1929, A50/B74). The striking creativity and freedom of human thought suggests that concepts are the sorts of mental states that can be freely recombined into novel structures, forming the finite basis of indefinitely many thinkable thoughts (e.g., Chomsky, 1986). The requirement of recombinability has roots in Evans' Generality Constraint: "if a subject can be credited with the thought that *a* is *F*, then he must have the conceptual resources for entertaining the thought that a is *G*, for every property of being *G* of which he has a conception" (1982, p. 104). The Generality Constraint is regularly taken as a constitutive condition on concept possession (Beck, 2012; Camp, 2009; Peacocke, 1992), sometimes in relaxed versions (Carruthers, 2004; 2009). Camp argues that STIMULUS-INDEPENDENCE captures the sense in which concepts are not merely recombinable, but recombinable in a way that constitutes "active, genuinely rational thinking" (Camp, 2009, p. 287).

2.2 | Pragmatism vs. representationalism

A starting point for Camp's approach is that concepts are mental abilities (Camp, 2009, p. 278n3; cp. Burge, 2010a, p. 197). A primary goal of her defense of STIMULUS-INDEPENDENCE is to develop a theory of concepts "that captures the core set of cognitive tasks that we expect concepts to perform" (Camp, 2009, p. 276). Camp's theory is thus a version of what Fodor calls "concept pragmatism" (2004, p. 30) and dubs "*the* characteristic doctrine of twentieth century philosophy of mind/language" (2004, p. 29; emphasis his). According to concept pragmatism, "concept possession is some sort of dispositional, epistemic condition" (*ibid.*). Camp's defense of STIMULUS-INDEPENDENCE commits to a form of concept pragmatism on which the ability to form novel thoughts independently of stimulation is constitutive of concept possession.[1]

Fodor's alternative to concept pragmatism is "Cartesianism", according to which to "have the concept DOG is to be able to think about dogs as such" (2004, p. 31). Cartesianism is a thesis about what it is to possess a concept, but it fits naturally with a representationalist theory of concepts themselves. That is, concepts are mental representations (i.e., particulars rather than abilities) and the "concept DOG is that mental particular the possession of which allows one to represent—to bring before one's mind—dogs as such" (Fodor, 2003, p. 19). Cartesianism is thus the natural extension of representationalism (i.e., "the representational theory of mind" (Fodor, 1998)), according to which mental states are primarily (relations to) representations, and mental processes are primarily computational operations over representations.

It is crucial to representationalism that mental representations are ontologically more basic than cognitive abilities. Cognitive abilities, such as the ability to draw inferences, are analyzed in terms of mental processes, such as inferences; and mental processes are in turn analyzed as operations over mental representations, such as the operations over constituent structures that underlie deductive inference (Fodor & Pylyshyn, 1988; Quilty-Dunn & Mandelbaum, 2018). If this representationalist story is correct, then it can't also turn out that mental representations are analyzed as abstractions over cognitive abilities, on pain of circularity. If our cognitive abilities arise out of processes defined over representations, then representations must be characterizable independently of those abilities. Representationalism provides just such an independent characterization: mental representations are symbols, i.e., vehicles with representational contents.

Concept pragmatism is motivated by a desire to capture a "core set of cognitive tasks" (Camp, 2009, p. 276) by building the ability to perform them into the metaphysics of concepts (cp. Prinz & Clark, 2004). For representationalists, cognitive abilities can be *data* to be explained by positing concepts, and some can even be *diagnostic* of concepts (rather than other sorts of representations, or nonrepresentational states). But they can't be *constitutive* of concepts or their possession conditions. A creature possesses cognitive abilities in virtue of having the right sort of computational machinery for computing over the right sort of concepts; possessing the concepts explains possessing the abilities rather than vice versa. Concept pragmatism, according to representationalism, conflates epistemology and metaphysics. While abilities to accomplish various cognitive tasks might be excellent evidence to justify attributions of concepts to some creature, the attributions are not made true by the creature's possessing the abilities.

This representationalist critique of concept pragmatism is, to be sure, controversial.[2] Providing a full defense of representationalism isn't possible here. However, one needn't be convinced to be interested in the upshots of a representationalist approach. One need only leave open the possibility that concepts are symbols (Camp, 2009, p. 278n3; cp. Evans, 1982, pp. 100–101). If concepts are symbols, then it's an open question whether inferential or other cognitive abilities are constitutive of concept possession. At first glance, tokening a symbol need not presuppose an ability to transform the symbol in any particular way, and thus tokening a concept need not require the ability to use it in stimulus-independent thought.

Furthermore, concept pragmatists nearly always grant that cognitive abilities don't constitute absolute possession conditions on concepts. In his original formulation of the Generality Constraint, Evans added in a footnote that there must be "a proviso about the categorial appropriateness of the predicates to the subjects" (1982, p. 101n17). Camp (2004) rejects Evans' proviso, but, following Peacocke (1992, pp. 42–43), she grants that "strange chemical reactions, psychological traumas, or other external factors" (Camp, 2009, pp. 278–279) may prevent recombination.

These pathological cases are not taken to undermine "the conceptual abilities themselves" (Camp, 2009, p. 279). However, barriers to the cognitive use of concepts in inference or recombination need not arise pathologically. As Peacocke notes, such barriers can arise "at the level of hardware" (1992, p. 43). There's no obvious reason why such hardware-level factors might not happen to be built-in to the normal functioning of some minds—that is, they might be aspects of *mental architecture*. Mental architecture comprises, roughly, functional properties that are invariant across changes in representational content, such as distinctions between memory stores (e.g., working memory vs. long-term memory) (Pylyshyn, 1984, pp. 30–32). There could in principle be aspects of mental architecture that prevent concepts from being deployed outside certain limited contexts.

One way to make sense of this possibility is to consider memory limitations. Imagine a thinker with extremely limited working-memory capacity. Suppose they possess the complex concept HORSE THAT IS SMALLER THAN THE LARGEST CLYDESDALE ON EARTH but, every time they try to compose that concept into a thought like SEABISCUIT'S SECOND-YOUNGEST OFFSPRING IS A HORSE THAT…, memory resources fail and the structure crashes before it's fully formed. This thinker lacks the ability freely to recombine this concept, but they possess the concept nonetheless; the limitation lies not in the concept itself, but in working-memory capacity. Thus a species might evolve that grasps and stores concepts but for *completely independent, non-pathological reasons* lacks the ability to deploy them in certain ways—due to working-memory limitations or other background architectural factors.

A similar possibility arises regarding stimulus-independence *tout court*. A creature might possess a concept that it uses for perceptual identification. However, sustaining its deployment in the absence of relevant sensory input requires working-memory resources. It thus seems metaphysically possible that a creature might have a concept and the ability to deploy it in response to stimulation while also possessing a mental architecture that precludes deploying the concept independently of stimulation.

One way to see the coherence of this possibility is to consider possible changes in mental architecture. Perhaps one day we'll be able to insert chips that enhance working-memory resources into brains. Suppose the following counterfactual is true: if we were to insert a chip into the brain of the creature just described and enhance their memory resources, they would be able to sustain the deployment of the concept independently of stimulation and freely recombine it with their other concepts. Perhaps the creature tokens representations that are apt to function as premises in *modus ponens* inferences but lacks the memory resources required to make those inferences. However, it might be true that they would gain the ability to make them were we to insert the memory-enhancing chip.

A pragmatist might balk at invoking such fanciful scenarios. But the deeper truth they illustrate is that stimulus-independent use in inference and recombination require more than the mere possession of a concept. They require the right background architectural setup as well. And if we are willing to grant "hardware-level" restrictions on stimulus-independent deployment, then there's no clear reason why we should deny that concept possession could survive such restrictions when they arise from the architecture rather than pathology.

According to Fodor's Cartesianism, possessing DOG requires only the ability to think about dogs as such. But the present line of reasoning suggests that even Fodor's view is too pragmatist. Suppose (as Fodor surely would) that concepts are symbols and that their use in thought depends on background features of mental architecture. In that case, thinking about dogs as such requires that a symbol be retrieved from memory and deployed. Memory retrieval, however, is a psychological process that can fail. Failures of retrieval are perfectly ordinary and don't entail that the relevant information fails to be stored, as the head-slaps that follow the revelation of an answer to a pub trivia question may attest. It's therefore possible that some symbol might be stored and yet not be retrievable for independent reasons (e.g., mechanisms underwriting retrieval have malfunctioned). The concept would remain stored nonetheless. While this scenario involves malfunction, it is again conceivable that the same factor could arise from non-pathological aspects of mental architecture. Suppose some creature evolved innate symbols, but over time those symbols lost their adaptiveness in the environments of that creature's descendants. It's conceivable that the vestigial symbols remain innately stored in the minds of those descendants, but subsequent evolution has rendered them irretrievable as a matter of course.

Thus if concepts are symbols, then having a concept need not presuppose *any* cognitive abilities—not even the ability to use the concept in thought. Instead, having a concept is simply *storing a certain symbol in memory*. This view—call it "possession-as-storage"—takes maximally seriously the idea that the mind is a computational system and that concepts are symbols stored and computed over in that system.

Possession-as-storage may strike many readers as extreme. However, it is worth clarifying the relevant notion of *cognitive ability*. A natural way of interpreting 'ability' in the concepts literature (and in this paper thus far) is in terms of an ability that a creature can exercise at a time; the question whether a concept is possessed by that creature is (for pragmatists) transformed into a question about what the creature can do at that time. In that sense of ability, possession-as-storage entails that no cognitive abilities are constitutive of concept possession.

One might more permissively attribute abilities to creatures who cannot exercise them in nearby possible worlds (e.g., because of architectural limitations). One might also attribute abilities not to creatures but to *concepts themselves*—Camp writes that external factors limiting recombinability fail to affect "the conceptual capacities themselves" (2009, p. 279; cp. Peacocke, 1992, p. 43). Combining these ideas, we might develop the following modified pragmatist view: possessing a concept requires that the creature possesses a state that has the ability to be used in stimulus-independent recombination and inference given the right circumstances, including background mental architecture.

Even a Fodorian representationalist could accept this modified pragmatist view. However, a primary motivation for pragmatism is that it ties concept possession to verifiable cognitive tasks and thus furnishes us with diagnostic tests of conceptuality emanating from a "practically useful account that captures the core set of cognitive tasks that we expect concepts to perform" (Camp, 2009, p. 276). The modified pragmatist view weakens the link between concept possession and the actual exercise of cognitive abilities in a way that vitiates the initial motivation behind concept pragmatism. It also allows—importantly for present purposes—that creatures that cannot actually exercise stimulus-independent cognitive abilities may possess concepts nonetheless.

What other evidence could license concept attribution beyond stimulus-independent cognition? For a representationalist, what matters is *what type of symbol* is instantiated in the mind rather than the cognitive abilities possessed by the creature. Some representations aren't concepts, such as icons, since they lack the right sort of representational format. Thus I suggest that investigating the representational format of perception provides an independent means of answering these questions about conceptuality. In particular, we can investigate whether some perceptual representations have a predicate-argument structure that is usable for logical inference (given the right background architecture). Such representations might be conceptual even if the creatures that possess them lack paradigmatically conceptual cognitive abilities.

3 | OBJECTS, PREDICATION, AND THE SYNTAX OF PERCEPTION

For both Burge and Camp, the argument for STIMULUS-INDEPENDENCE relies on a more fundamental aspect of concepts: compositionality. Concepts can compose into more complex structures, as when PET and FISH combine in PET FISH. In particular, concepts compose into truth-evaluable propositional structures, like THIS IS A FISH. What marks the simplest propositional structures is their *predicate-argument structure*. A picture of a fish might represent an object and even represent it as a fish, but it doesn't do so by means of a predicate-argument structure.

It's important to distinguish predication as an aspect of *content* (a structural feature of a proposition) and as an aspect of *format*. One might argue that a picture expresses predication in that its content predicates a property of some individual, but predicate-argument structure is not explicit in the structure of the vehicle (i.e., its format). Minimally, predicate-argument structure requires that "some sort of functional relation among syntactic constituents maps onto some sort of logical or metaphysical relation among the semantic values of those constituents" (Camp, 2007, p. 157). In a sentence like 'This is a fish', "the syntactic relation of function application mirrors a metaphysical relation of instantiation" (*ibid*.); the constituent 'This' corresponds to the individual, 'fish' corresponds to the property *fish*, and the syntactic relation between them functions to express the instantiation of *fish* by the individual. This sort of structure is a canonical example of predicate-argument structure (where 'fish' functions as predicate and 'This' as argument).

In a picture of a fish, this structure is absent. There are not two separate constituents standing for an individual and for the property of being a fish. Instead, the same part of the picture that represents the individual also represents its various properties. In this sense, iconic representations are "holistic" (Green & Quilty-Dunn, 2017; Quilty-Dunn, 2019a). They're not *digital*, in Camp's sense of taking "a small number (typically, a singleton or pair) of discrete elements as inputs" (Camp, 2018, p. 25). Icons have a comparatively large number of primitives (e.g., pixels—cf. Davies, 2020), and their primitives encode multiple semantic values at once. For example, a part of a picture might encode values along multiple spatial axes as well as features instantiated at the corresponding location, such as values along color dimensions, shape and size dimensions, etc. A depicted individual is represented by means of parts (primitives or regions) of the icon that encode other information as well, including parts of the individual and/or their values along spatiotemporal and featural dimensions (Hagueland, 1998, p. 192; Kulvicki, 2006, p. 125; ms).

I've argued elsewhere that perceptual object representations ("PORs"), the representations we use to perceptually detect and track objects, have a discursive/digital rather than iconic/analog format (Green & Quilty-Dunn, 2017; Quilty-Dunn, 2019b). I'll briefly describe these arguments now.

First, PORs comprise separate constituents for *individuals and properties*. Tracking via PORs involves an index-like constituent that picks out individual objects and continues to track them even when featural information changes (Zhou et al., 2010) or is lost altogether (Bahrami, 2003; Kibbe & Leslie, 2011; Scholl, Pylyshyn, & Franconeri, 1999). The best and simplest explanation of these tracking abilities posits discrete constituents for individuals that are non-holistically bound to featural information (Pylyshyn, 2003; Scholl & Leslie, 1999), which fits a discursive model better than an iconic one.

Second, PORs comprise separate constituents for *distinct feature dimensions*. While icons represent (e.g.) the color and orientation of a triangle by means of the same parts of the icon, PORs can successfully encode both features but lose them independently of each other in visual short-term memory (Bays, Wu, & Husain, 2011; Dowd & Golomb, 2019; Fougnie & Alvarez, 2011; Fougnie, Cormeia, & Alvarez, 2013; Markov, Tiurina, & Utochkin, 2019; Wang, Cao, Theeuwes, Olivers, & Wang, 2017; Markov et al. ms). The separability of features in PORs suggests that distinct features are represented *via* distinct vehicles, implicating discursive format.

Third, PORs comprise separate constituents for *high-level vs. low-level features*. Experimental evidence concerning storage of property-information in PORs shows that they represent properties like *hammer* independently of other features (Goodhew, Greenwood, & Edwards, 2016; Gordon, 2014; Gordon & Vollmer, 2010; Gordon, Vollmer, & Frankl, 2008; Pollatsek, Rayner, &

Collins, 1984; 1990). For example, previewing the word ‘hammer’ in a perceptual object yields an object-specific benefit in discriminating the visual appearance of a hammer (Gordon & Irwin, 2000), and a previewed image of a hammer yields an object-specific benefit in discriminating the sound of a hammer from that object’s location (Jordan, Clark, & Mitroff, 2010). The POR therefore encodes *hammer* in a way that is not tied to any low-level features, even generic ones (cf. Burge, 2014). These effects are predicted and explained by the hypothesis that high-level properties are represented *via* discrete constituents in PORs rather than *via* icons.

The best explanation of these effects is that PORs have a discursive representational format (Quilty-Dunn, 2019b). What’s most relevant for present purposes is the particular kind of discursive format such explanations appeal to. Arguably, maps are at least partly discursive in virtue of symbols that function as markers (Blumson, 2012; Camp, 2007; Kulvicki, 2015b). But the format of maps doesn’t involve predicate-argument structure as a sentence does, even if, as some argue, the *content* of maps is predicative (Casati & Varzi, 1999; Kulvicki, 2015b; cf. Camp, 2018; Rescorla, 2009b; 2009c).

In the case of PORs, however, individuals are represented by discrete constituents, and separate constituents non-holistically represent various high- and low-level properties. The function of PORs is to encode information about individual objects and bind features by attributing them to the same individual. In that case, it seems true that “some sort of functional relation among syntactic constituents maps onto some sort of logical or metaphysical relation among the semantic values of those constituents” (Camp, 2007, p. 157). The constituents of a POR that represent properties are functionally related to the constituent that represents the object in a way that makes the POR accurate iff the represented object instantiates (or recently instantiated) the represented properties.

What about Burge’s notion of “pure predication”? As articulated by Burge, pure predication seems to presuppose stimulus-independent abilities and is thus ill-suited to apply to perception. However, perhaps pure predication can be understood without presupposing stimulus-independence. For Burge, pure predication requires that an attributive function as a main predicate outside the scope of a referential noun-phrase-like structure: e.g., ‘That F is G’ contains a main predicate ‘G’ and an attributive ‘F’ that functions to guide reference of ‘That F’. Since Burge claims that in perception “this attributive function always serves and is subordinate to the larger perceptual function of identificational reference” (2010b, p. 41), he concludes that perception cannot incorporate pure predication.

Burge’s claim that perceptual attributives invariably guide reference is an empirical one. PORs provide counterexamples. To track moving objects, the visual system must determine whether an object at a previously unoccupied location is a new object, or a previously perceived object that moved—this is the *correspondence problem* (Kahneman, Treisman, & Gibbs, 1992). Solving the correspondence problem requires using information about an object to guide continuous visual reference to it. We can therefore test the claim that all perceptual attributives guide perceptual reference by examining the role of attributives in object correspondence.

Object correspondence has long been thought to rely on spatiotemporal information rather than other attributives in PORs (Flombaum, Scholl, & Santos, 2009; Kahneman et al., 1992). Recent evidence suggests that surface features like color can figure in object correspondence, which is consistent with Burge’s idea that attributives play a reference-guiding role (Moore, Stephens, & Hein, 2020). For example, changes in surface features like color can disrupt object correspondence (Jiang, 2020). Remarkably, however, whether a feature is usable for correspondence shifts depending on scene context (Quilty-Dunn & Green, unpublished).

Gordon et al. (2008) found that (i) changes in orientation disrupt object correspondence when objects are presented close to where subjects are visually fixating, but (ii) the same changes fail to disrupt correspondence when the objects appear closer to the periphery. One might object that perhaps orientation isn't encoded in the latter case, but Gordon et al. found that subjects accurately identified orientation changes, showing that PORs did store orientation information. In other words, when properties like orientation are encoded near foveal vision with high resolution, the visual system uses them to guide visual reference; however, when the object is further out and resolution is degraded, *the POR still attributes the property but it no longer guides visual reference*. The visual system works out that, even though the object has this property, poor representational quality means it should not guide reference. This result suggests a syntactic movement of the attributive facing leftward from within the scope of the noun phrase That object to the main predicate position, constituting pure predication of the form That object is facing leftward.

Furthermore, at least some of the constituents of PORs behave like ordinary predicative concepts that denote categories like *hammer*. Explaining results such as Gordon and Irwin's (1996; 2000) and Jordan et al. (2010)) (discussed above) requires positing discursive representations of categories like *hammer* in PORs that (i) represent kinds independently of low-level features, (ii) transcend individual sensory modalities, and (iii) are mapped to lexical items. These features are characteristic of concepts like HAMMER. A salient alternative explanation is that the initial previewed feature (e.g., 'hammer') simply triggers associations with other features such as related wordforms, images, and sounds. In that case, wordforms like 'hammer' should activate associated wordforms like 'nail'. But while those associations are triggered, they do not yield any *object-specific* benefit (Gordon & Irwin, 1996). Instead, the object-specific benefit is limited to information that falls under the concept HAMMER (e.g., the sound and appearance of hammers), just as would be predicted if HAMMER were a constituent of the POR.

These considerations in favor of conceptual constituents of PORs—obeying hallmarks of conceptuality and exhibiting a syntactic distinction between main predicates and reference-guiding attributives—intersect nicely in other recent object-correspondence experiments. Specifically, it appears that which low-level attributives guide reference on a particular occasion is driven by (*inter alia*) which conceptual predicates figure in PORs.

Color is often ignored for object correspondence unless scene context renders it a useful cue to object continuity (Papenmeier, Meyerhoff, Jahn, & Huff, 2014). Gordon and Vollmer (2010) used easily categorizable stimuli (e.g., buckets and bananas) and tested whether color changes disrupted object correspondence. They found a divergence: color changes did disrupt correspondence, but only for objects whose categories have diagnostic colors. Yellow is diagnostic of bananas, while buckets can be any color you like. Color changes correspondingly disrupted object correspondence for bananas but not buckets. Whether color guides reference is thus sensitive to color's usefulness as a cue not only to the particular (token) object in the scene, but also its usefulness for discriminating object categories.

Even more striking results come from a study on "Ternus" apparent-motion displays, which involve a serial presentation of objects in a way that's ambiguous between (i) the leftmost object hopping over the other(s) ("element motion") and (ii) all objects shifting rightward ("group motion") (Figure 1). Apparent-motion displays in general (Odic, Libertus, Feigenson, & Halberda, 2013), and Ternus displays in particular (Stepper, Moore, Rolke, & Hein, 2019), typically probe object-correspondence processes.

Hsu, Taylor, and Pratt (2015) used frogs as stimuli. Since frogs hop forward, they tested whether forward-facing frogs biased correspondence toward element motion, i.e., seeing the leftmost frog "hop" over the other. They found that, indeed, subjects were more likely to see element motion

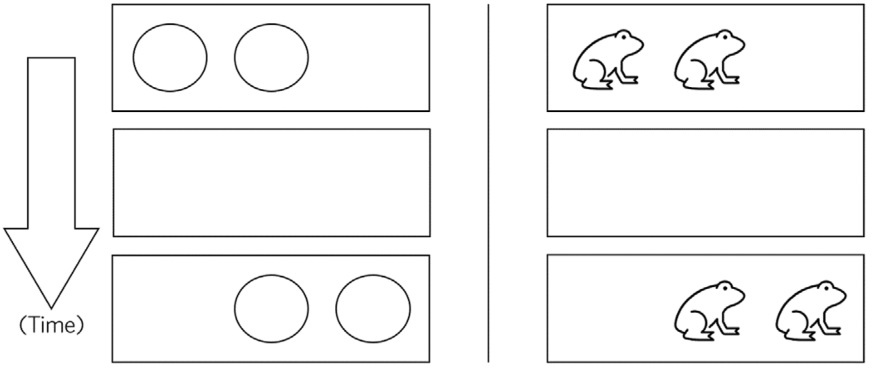

FIGURE 1 Ternus displays with circles (left) and frogs (right)

than group motion for forward-facing vs. backward-facing frogs. This result provides evidence that concepts like FROG can figure in PORs and determine whether orientation guides perceptual reference.

Perhaps the effect is due to (a) orientation *tout court* rather than anything frog-specific, or (b) an association between frog-like visual attributives and forward motion. Hsu et al. controlled for these possibilities, however, finding that (a) there was no difference between forward vs. backward-oriented triangles, and that (b) the difference *did* persist for visually unrealistic line drawings of frogs, suggesting that the effect is not due to low-level visual factors. Instead, it seems that whether orientation guides object-based reference is controlled partly by the conceptual category of the stimulus and whether it indicates that orientation is useful in resolving object correspondence. Concepts thus not only figure in PORs as predicates, they also shape the role of PORs in core perceptual functioning and determine the syntactic position of other attributives like color and orientation as either reference-guiding attributives or main predicates.

The version of conceptualism defended here requires a delicate balance. Since PORs include concepts organized in a predicate-argument structure, they must be apt to function in logical inference given the right background cognitive architecture. But if we reject STIMULUS-INDEPENDENCE and concept pragmatism, possessing conceptualized PORs can't require possessing logical-inferential abilities. Thus conceptualism about PORs makes some empirical predictions: we should be able to find PORs with a predicate-argument discursive format independently of logical-inferential abilities; but where logical-inferential abilities are present, we should find evidence that PORs can function as premises.

Evaluating these predictions requires investigating infants and non-human animals. Fortunately, there is significant evidence from developmental and comparative psychology concerning PORs. Object-tracking abilities arise in the first three months of life, and researchers generally agree that the same PORs are responsible for object-based effects in infants and adults (Carey, 2009; Carey & Xu, 2001; Scholl & Leslie, 1999). Multiple-object tracking has also been demonstrated in macaques (Anderson, Mitchell, & Reynolds, 2011; Mitchell, Sundberg, & Reynolds, 2007). While there's not much comparative evidence using traditional vision-science tests of PORs in animals (cf. Flombaum, Kundey, Santons, & Scholl, 2004), it is possible to look for evidence of PORs through their role in numerical cognition.

There are two ways we quickly enumerate sets of items without counting: *subitizing*, in which the precise numerosity of a small set is represented, and the *approximate number system*, which represents the approximate numerosity of indefinitely many items and discriminates sets by

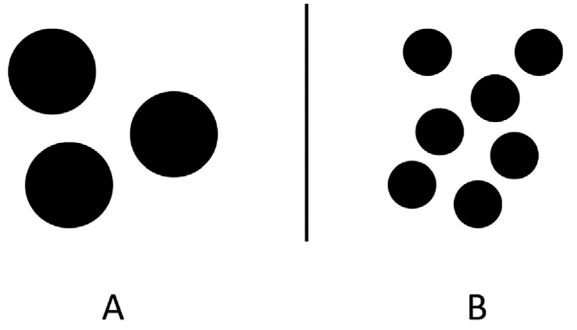

FIGURE 2 Subitizing vs. approximating

ratios rather than absolute cardinality, in accordance with Weber's Law (Carey, 2009; Feigenson, Dehaene, & Spelke, 2004). For instance, in Figure 2, most people can immediately see that A contains three items (subitizing); it's harder to say without counting how many items B contains, but we can easily see that B contains more items than A (approximate number system).

Due to its presence in infants and a characteristic set-size limit of 3 or 4, subitizing has long been thought to rely on PORs (Carey, 2009; Uller, Carey, Huntley-Fenner, & Klatt, 1999). In a striking confirmation of this hypothesis, Chesney and Haladjian (2011) had subjects perform a multiple-object tracking task while periodically subitizing a separate set of objects. They found that the maximum set size that could be subitized decreased by one for every item being tracked. Thus the very PORs used to track objects are recruited by numerical cognition to enumerate small sets. Moreover, while the approximate number system obeys Weber's Law, arguably a diagnostic feature of analog format (Beck, 2019; Clarke, forthcoming), subitizing does not (Choo & Franconeri, 2014).

Given these facts about subitizing in humans, we can look to the comparative numerical cognition literature to (dis)confirm the hypothesis that nonhuman animals possess PORs. Since PORs in humans are discursive and fulfill characteristic functions like tracking and subitizing, evidence of PORs that fulfill similar functions in other animals provides nondemonstrative but probative evidence for discursive format in animal minds.

Subitizing is present in many nonhuman animals, including birds (Rugani, Regolin, & Vallortigara, 2008) and horses (Uller & Lewis, 2009). Even animals typically thought to lack cognitive-inferential abilities can subitize. For example, guppies are sensitive to size differences in shoals of other guppies, and recent evidence shows they distinguish groups of 3-vs.-4 but not 6-vs.-8, suggesting that in addition to a ratio-sensitive approximate number system, they also have a POR-based subitizing system (Agrillo, Piffer, Bisazza, & Butterworth, 2012; Bisazza, Agrillo, & Lucon-Xiccato, 2014). Remarkably, this ability may exist in invertebrates: jumping spiders habituated to a display containing some prey will hesitate if, after approaching the display later, the number of prey has changed; while they are not sensitive to differences of 3-vs.-6, they are sensitive to 1-vs.-2 and 2-vs.-3, implicating POR-based numerical cognition (Cross & Jackson, 2017). Similar results have been shown for honeybees (Gross et al., 2009), and bumblebees recognize objects across perceptual modalities (Solvi, Al-Khudhairy, & Chittka, 2020).

I've argued that PORs are discursive and present in infants and animals. However, the evidence for discursive predicate-argument structure has thus far come from adult humans. It's conceivable that a different representational kind is responsible for POR-like effects in infants and animals. My response to this worry is twofold. First, while it's tricky to find evidence for instances of the same representational kind across species, a standard methodology is to look for "characteristic psychological effects" (Smortchkova & Murez, forthcoming), or "signatures" that are present "across developmental or evolutionary time" (Carey, 2009, p. 70). The signatures of PORs include use in tracking and subitizing; thus the evidence just reviewed should be taken as strong evidence for PORs in infants and animals.

Second, there's some encouraging recent evidence implicating discursive predicate-argument structure in infant PORs. Infants are known to keep track of the number and location of objects hidden behind occluders but fail to notice when a hidden object changes from a duck to a truck (Xu & Carey, 1996) until about 12 months (Xu et al., 2004). While this evidence may support discursive PORs with argument-terms (i.e., visual indexes) coming apart from featural and categorical predicates, one could reasonably object that relevant featural/categorical information is not bound to the infants' representations at all.

Kaldy and Leslie (2005) habituated six-month-old infants to displays in which two objects with different shapes (circle and triangle) were moved backward one at a time. During the experimental phase, each object ended up behind an occluder; then, when the occluders were later removed, the object may have changed its shape or not. They found that infants looked longer at the display when the object changed its shape, a signature of expectation violation. Interestingly, however, they only did so when the object was the second object hidden. For the first-hidden object, the POR degraded in memory and lost featural information. Kibbe and Leslie (2011) replicated this effect but added a third condition: the first-hidden object either persisted, changed shape, or vanished altogether. They found that, while infants lost shape information, they did look longer in the "vanish" condition. This result strongly suggests that these six-month-olds have PORs that index objects, track their location, and successfully encode featural information, but that the featural information can be lost while the index-like element remains, suggesting that "infants' working memory supports an object representation that is featureless" (Kibbe & Leslie, 2011, p. 1505). Thus a completely different experimental paradigm in infants independently substantiates the tracking data in adults that suggests that PORs have discrete constituents for objects and their properties, just as a conceptualist model of PORs predicts.

Recall that another form of evidence for conceptual PORs above concerned apparently conceptual categories predicated of objects in a way that comes apart from low-level features. Similar results in pre-linguistic infants would provide powerful independent evidence for predicate-argument structure in PORs while also bolstering the continuity of PORs across adult humans and creatures that lack language and higher cognition. Kibbe and Leslie (2019) used the same paradigm to probe six-month-olds' PORs of occluded objects, but instead of a circle and triangle the objects were a ball and a human head. If the first-hidden object was swapped with a featurally distinct object in the same category—e.g., striped ball vs. polka-dotted ball of different colors—infants failed to notice. However, if the objects are swapped *across conceptual categories*—e.g., head vs. ball—then infants did notice. This is precisely the result we would predict if we thought concepts like HUMAN could function as predicates in PORs in pre-linguistic creatures without being holistically bound to low-level features.

The foregoing evidence suggests that PORs with discursive predicate-argument structure are present in infants and animals. The previous section argued that concept possession doesn't presuppose stimulus-independent thought such as logical inference; however, representations with predicate-argument structure are apt to figure as premises in logical inference in creatures capable of such inferences. Assuming that PORs have predicate-argument structure, therefore, an immediate question is whether they emerge independently of logical-inferential abilities and figure in inference when such abilities arise.

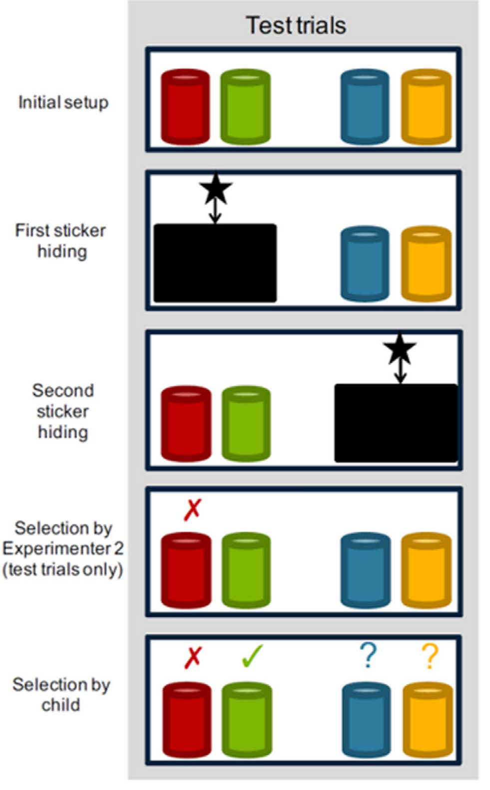

FIGURE 3 Mody and Carey's (2016) four-cup task (with permission from Elsevier) [Color figure can be viewed at wileyonlinelibrary.com]

While it is difficult to test for *modus ponens* inference without verbal stimuli (Burge, 2010b, p. 59), one can more easily test for other inferential patterns like disjunctive syllogism, i.e., *p-or-q; not-p; therefore, q*. Burge (2010b) in particular regards this pattern as a key diagnostic marker of inference. Suppose a dog chases a prey to a forked road, sniffs the first road and finds no scent, and immediately runs down the second road (Rescorla, 2009a). One might think that the dog reasons, *A-or-B; not-A; therefore, B* (Call, 2004). However, other architectures could generate this behavior without predicate-argument structure (Burge, 2010b; Rescorla, 2009a). For example, a creature could use a mental chalkboard to write two possibilities (*A; B*) and then erase *A* leaving only *B*, simulating disjunctive syllogism without actually performing it. How could we test whether creatures are using this simpler mental-chalkboard architecture or genuinely logical inference?

Mody and Carey (2016) devised a scenario where the two architectures come apart: the *four-cup task* (Figure 3). Children saw two stickers and two pairs of cups; each sticker went behind an occluder in front of one pair of cups. At this point, a creature capable of disjunctive syllogism could form representations *A-or-B* and *C-or-D*. A creature with just the mental-chalkboard architecture could only represent *A*; *B*; *C*; *D*, without any disjunction operator binding subsets of possibilities together. To test which architecture was present, Mody and Carey then showed subjects that cup *A* was empty. At this point, the two architectures license critically different updating procedures. The logical-inferential architecture represents *not-A*, which inferentially integrates with *A-or-B* to generate *B*. The mental-chalkboard architecture simply eliminates *A*, leaving *B*; *C*; *D*. Thus if children perform genuinely logical inferences, once they see *A* is empty they should immediately choose *B*; otherwise, they should be just as likely to go toward any of the three other cups.

They found that children fail the four-cup task until about age three, after which they successfully go for cup *B* more than chance. Since 2.5-year-olds fail to perform disjunctive syllogism but have long since developed PORs, this study provides evidence that discursive PORs predate logical-inferential abilities.

However, Mody and Carey suggest that the difference between 2.5-year-olds and 3-year-olds in this task coincides with a difference in linguistic capacities (2016, p. 47). One might conclude that logical-inferential abilities in older children and adults rely on language-mediated representations, not PORs. In a surprising follow-up study, however, Pepperberg, Gray, Cornero, Mody, and Carey (2019) found that Griffin, a Grey parrot, succeeds at the four-cup task. Thus the logical-inferential abilities required to perform disjunctive syllogism seem not to rely on language-mediated representations. Instead, Griffin's superior visual-working-memory capacity, or some such background architectural factor, allows him to perform inferences unavailable to 2.5-year-olds (Pepperberg et al., 2019, pp. 27ff; Porot, 2019).

Pepperberg et al. take Griffin's success to rely on object-based representations in visual working memory, which are plausibly PORs (Gao, Gao, Li, Sun, & Shen, 2011; Quilty-Dunn, 2019a; 2019b). Though Pepperberg et al. deny that Griffin's representations are propositional (2019, p. 25), their argument depends on the assumption that PORs are iconic.[3] If we reject that assumption, then the following hypothesis becomes a live possibility: PORs are discursive representations with predicate-argument structure—that is, conceptualized propositional structures—that arise independently of logical-inferential abilities and, once such abilities develop, can figure as premises in logical inferences. This hypothesis is parsimonious and consistent with both the available evidence and a representationalist view of concepts.

4 | CONCLUSION

This paper has provided a sketch of one form conceptualism might take that is grounded in perceptual, developmental, and comparative psychology. On this sketch, our most primitive capacities to track and enumerate objects employ conceptualized propositional structures. These structures later come to constitute a form of perceptual evidence on the basis of which we draw inferences about the perceptible world. Conceptuality is not constitutively dependent on stimulus-independent cognitive abilities, but rather provides a representational basis out of which those abilities grow.[4]

ENDNOTES

1 I focus here on cognitive *abilities* rather than capacities (Schellenberg, 2018).

2 Prinz and Clark, for example, argue that Fodor's representationalism threatens to "sever the apparent links between concept-having and any kinds of abilities or dispositions to act" (2004, p. 65). *Pace* Prinz and Clark, denying that concept possession is *constituted by* abilities doesn't require severing links from the former to the latter. Feet are not constituted by the ability to wear shoes, but accepting that fact doesn't sever the apparent links between having feet and having the ability to wear shoes.

3 Rescorla's (2009a) explanation in terms of probabilistic updating over cognitive maps also remains live. But Mody and Carey found that the children who succeeded at the four-cup task were *no less likely* to pick the right cup in that task than in a task where they saw the sticker go behind an occluder that contained a single bucket. This suggests that in the four-cup task, they aren't merely bumping up credences but rather exhibiting the "all-or-nothing character of deductive inference" (Burge, 2010b, p. 64).

[4] I'm grateful to audiences at the Varieties of Intentionality workshop at the New School in May 2019 and the Salzburg Mental Files Workshop in June 2020, and to Zed Adams, Sam Clarke, Alex Grzankowski, Zoe Jenkin, Chris Hill, Eric Mandelbaum, Michael Martin, David Papineau, Christopher Peacocke, Nick Shea, Wayne Wu, and especially Liz Camp for comments/discussion. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 681422.

REFERENCES

Agrillo, C., Piffer, L., Bisazza, A., & Butterworth, B. (2012). Evidence for two numerical systems that are similar in humans and guppies. *PLoS One*, *7*(2), e31923. https://doi.org/10.1371/journal.pone.0031923

Anderson, E. B., Mitchell, J. F., & Reynolds, J. H. (2011). Attentional modulation of firing rate varies with burstiness across putative pyramidal neurons in macaque visual area V4. *Journal of Neuroscience*, *31*(30), 10983–10992. https://doi.org/10.1523/jneurosci.0027-11.2011

Bahrami, B. (2003). Object property encoding and change blindness in multiple object tracking. *Visual Cognition*, *10*(8), 949–963. https://doi.org/10.1080/13506280344000158

Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, *49*(6), 1622–1631. https://doi.org/10.1016/j.neuropsychologia.2010.12.023

Beck, J. (2012). The generality constraint and the structure of thought. *Mind*, *121*(483), 563–600. https://doi.org/10.1093/mind/fzs077

Beck, J. (2018). Marking the perception-cognition boundary: The criterion of stimulus dependence. *Australasian Journal of Philosophy*, *96*(2), 319–334. https://doi.org/10.1080/00048402.2017.1329329

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

Bengson, J., Grube, E., & Korman, D. (2011). A new framework for conceptualism. *Noûs*, *45*(1), 167–189. https://doi.org/10.1111/j.1468-0068.2010.00764.x

Bermudez, J. L. (1998). *The paradox of self-consciousness*. Cambridge, MA: MIT Press.

Bisazza, A., Agrillo, C., & Lucon-Xiccato, T. (2014). Extensive training extends numerical abilities of guppies. *Animal Cognition*, *17*(6), 1413–1419. https://doi.org/10.1007/s10071-014-0759-7

Block, N. (manuscript). *The border between seeing and thinking*.

Blumson, B. (2012). Mental maps. *Philosophy and Phenomenological Research*, *85*(2), 413–434. https://doi.org/10.1111/j.1933-1592.2011.00499.x

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

Burge, T. (2010a). *Origins of objectivity*. Oxford: OUP.

Burge, T. (2010b). Steps toward the origin of propositional thought. *Disputatio*, *4*(29), 39–67. https://doi.org/10.2478/disp-2010-0010

Burge, T. (2014). Reply to Block: Adaptation and the upper border of perception. *Philosophy and Phenomenological Research*, *89*(3), 573–583. https://doi.org/10.1111/phpr.12136

Burnston, D. C. (2017a). Cognitive penetration and the cognition–perception interface. *Synthese*, *194*, 3645–3688. https://doi.org/10.1007/s11229-016-1116-y

Burnston, D.C. (2017b). Interface problems in the explanation of action. *Philosophical Explorations*, *20*(2), 242–258. https://doi.org/10.1080/13869795.2017.1312504

Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, *88*(1), 119–145. https://doi.org/10.1111/j.1933-1592.2012.00604.x

Call, J. (2004). Inferences about the location of food in the Great Apes (*Pan paniscus*, *Pan troglodytes*, *Gorilla gorilla*, and *Pongo pygmaeus*). *Journal of Comparative Psychology*, *118*(2), 232–241. https://doi.org/10.1037/0735-7036.118.2.232

Camp, E. (2004). The generality constraint and categorial restrictions. *The Philosophical Quarterly*, *54*(215), 209–231. https://doi.org/10.1111/j.0031-8094.2004.00348.x

Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, *21*, 145–182. https://doi.org/10.1111/j.1520-8583.2007.00124.x

Camp, E. (2009). Putting thoughts to work: Concepts, systematicity, and stimulus-independence. *Philosophy and Phenomenological Research*, *78*(2), 275–311. https://doi.org/10.1111/j.1933-1592.2009.00245.x

Camp, E. (2018). Why maps are not propositional. In A. Grzankowski & M. Montague (Eds.), *Non-propositional intentionality* (pp. 19–45). Oxford: OUP.).

Carey, S. (2009). *The origin of concepts*. Oxford: OUP.

Carey, S., & Xu, F. (2001). Infants' knowledge of objects: Beyond object files and object tracking. *Cognition*, *80*, 179–213. https://doi.org/10.1016/s0010-0277(00)00154-2

Carruthers, P. (2004). On being simple minded. *American Philosophical Quarterly*, *41*(3), 205–220. https://doi.org/10.1093/0199277362.003.0012

Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 89–107). Cambridge: CUP.

Casati, R., & Varzi, A. C. (1999). *Parts and places: The structures of spatial representation*. Cambridge, MA: MIT Press.

Chesney, D. L., & Haladjian, H. H. (2011). Evidence for a shared mechanism used in multiple-object tracking and subitizing. *Attention, Perception, & Psychophysics*, *73*, 2457–2480. https://doi.org/10.3758/s13414-011-0204-9

Chomsky, N. (1986). *Knowledge of language: Its nature, origin, and use*. Westport, CT: Praeger.

Choo, H., & Franconeri, S. L. (2014). Enumeration of small collections violates Weber's Law. *Psychonomic Bulletin and Review*, *21*, 93–99. https://doi.org/10.3758/s13423-013-0474-4

Clarke, S. (forthcoming). Beyond the icon: Core cognition and the bounds of perception. *Mind & Language*, https://doi.org/10.1111/mila.12315

Cross, F. R., & Jackson, R. R. (2017). Representation of different exact numbers of prey by a spider-eating predator. *Interface Focus*, *7*(3), 20160035. https://doi.org/10.1098/rsfs.2016.0035

Davies, W. (2020). Colour relations in form. *Philosophy and Phenomenological Research*, https://doi.org/10.1111/phpr.12679

Dowd, E. W., & Golomb, J. D. (2019). Object-feature binding survives dynamic shifts of attention. *Psychological Science*, *30*(3), 343–361. https://doi.org/10.1177/0956797618818481

Evans, G. (1982). *The varieties of reference*. Oxford: OUP.

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, *8*(7), 307–314. https://doi.org/10.1016/j.tics.2004.05.002

Flombaum, J. I., Kundey, S. M., Santons, L. R., & Scholl, B. J. (2004). Dynamic object individuation in rhesus macaques: A study of the tunnel effect. *Psychological Science*, *15*(12), 795–800. https://doi.org/10.1111/j.0956-7976.2004.00758.x

Flombaum, J. I., Scholl, B. J., & Santos, L. R. (2009). Spatiotemporal priority as a fundamental principle of object persistence. In B. Hood & L. R. Santos (Eds.), *The origins of object knowledge*. Oxford: OUP.

Fodor, J. A. (1983). *The modularity of mind*. MIT press.

Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. Oxford: Clarendon.

Fodor, J. A. (2003). *Hume variations*. New York: OUP.

Fodor, J. A. (2004). Having concepts: A brief refutation of the twentieth century. *Mind and Language*, *19*(1), 29–47. https://doi.org/10.1111/j.1468-0017.2004.00245.x

Fodor, J. A. (2007). The revenge of the given. In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 105–116). Oxford: Blackwell.

Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71. https://doi.org/10.7551/mitpress/2103.003.0002

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature–store model. *Journal of Vision*, *11*(12), 1–12. https://doi.org/10.1167/11.12.3

Fougnie, D., Cormeia, S. M., & Alvarez, G. A. (2013). Object-based benefits without object-based representations. *Journal of Experimental Psychology: General*, *142*(3), 621–626. https://doi.org/10.1167/12.9.300

Gao, T., Gao, Z., Li, J., Sun, Z., & Shen, M. (2011). The perceptual root of object-based storage: An interactive model of perception and visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(6), 1803–1823. https://doi.org/10.1037/a0025637

Goodhew, S. C., Greenwood, J. A., & Edwards, E. (2016). Categorical information influences conscious perception: An interaction between object-substitution masking and repetition blindness. *Attention, Perception, & Psychophysics*, *78*, 1186–1202. https://doi.org/10.3758/s13414-016-1073-z

Gordon, R. D. (2014). Saccade latency reveals episodic representation of object color. *Attention, Perception, & Psychophysics*, *76*(6), 1765–1777. https://doi.org/10.3758/s13414-014-0683-6

Gordon, R. D., & Irwin, D. E. (1996). What's in an object file? Evidence from priming studies. *Perception and Psychophysics*, *58*(8), 1260–1277. https://doi.org/10.3758/bf03207558

Gordon, R. D., & Irwin, D. E. (2000). The role of physical and conceptual properties in preserving object continuity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(1), 136–150. https://doi.org/10.1037/0278-7393.26.1.136

Gordon, R. D., & Vollmer, S. D. (2010). Episodic representation of diagnostic and non-diagnostic object color. *Visual Cognition*, *18*(5), 728–750. https://doi.org/10.1080/13506280903004190

Gordon, R. D., Vollmer, S. D., & Frankl, M. L. (2008). Object continuity and the transsaccadic representation of form. *Perception and Psychophysics*, *70*, 667–679. https://doi.org/10.3758/pp.70.4.667

Green, E. J., & Quilty-Dunn, J. (2017). What is an object file? *British Journal for the Philosophy of Science*, https://doi.org/10.1093/bjps/axx055

Gross, H. J., Pahl, M., Si, A., Zhu, H., Tautz, J., & Zhang, S. (2009). Number-based visual generalization in the honeybee. *PLoS One*, *4*(1), e4263. https://doi.org/10.1371/journal.pone.0004263

Haugeland, J. (1998). Representational genera. In his *Having thought* (pp. 171–206). Cambridge, MA: Harvard University Press.

Hopp, W. (2011). *Perception and knowledge: A phenomenological account*. Cambridge University Press.

Hsu, P., Taylor, J. E. T., & Pratt, J. (2015). Frogs jump forward: Semantic knowledge influences the perception of element motion in the Ternus display. *Perception*, *44*(7), 779–789. https://doi.org/10.1177/0301006615596903

Jacobson, H., & Putnam, H. (2016). Against perceptual conceptualism. *International Journal of Philosophical Studies*, *24*(1), 1–25. https://doi.org/10.1080/09672559.2015.1047164

Jiang, H. (2020). Effects of transient and nontransient changes of surface feature on object correspondence. *Perception*, *49*(4), 452–467. https://doi.org/10.1177/0301006620913238

Jordan, K. E., Clark, K., & Mitroff, S. M. (2010). See an object, hear an object file: Object correspondence transcends sensory modality. *Visual Cognition*, *18*(4), 492–503. https://doi.org/10.1080/13506280903338911

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, *24*(2), 175–219. https://doi.org/10.1016/0010-0285(92)90007-o

Kaldy, Z., & Leslie, A. M. (2005). A memory span of one? Object identification in 6.5-month-old infants. *Cognition*, *97*, 153–177. https://doi.org/10.1016/j.cognition.2004.09.009

Kant, I. (1929). *Critique of pure reason*. N. Kemp Smith (Trans.). London: Macmillan. (Original work published 1781).

Kibbe, M. M., & Leslie, A. M. (2011). What do infants remember when they forget? Location and identity in 6-month-olds' memory for objects. *Psychological Science*, *22*(12), 1500–1505. https://doi.org/10.1177/0956797611420165

Kibbe, M. M., & Leslie, A. M. (2019). Conceptually rich, perceptually sparse: Object representations in 6-month-old infants' working memory. *Psychological Science*, https://doi.org/10.1177/0956797618817754

Kulvicki, J. (2006). *On images*. Oxford: OUP.

Kulvicki, J. (2015a). Analog representation and the parts principle. *Review of Philosophy and Psychology*, *6*(1), 165–180. https://doi.org/10.1007/s13164-014-0218-z

Kulvicki, J. (2015b). Maps, pictures, and predication. *Ergo*, *2*(7). https://doi.org/10.3998/ergo.12405314.0002.007

Mandelbaum, E. (2018). Seeing and conceptualizing: Modularity and the shallow contents of perception. *Philosophy and Phenomenological Research*, *97*(2), 267–283. https://doi.org/10.1111/phpr.12368

Mandik, P. (2012). Color-consciousness conceptualism. *Consciousness and Cognition*, *21*, 617–631. https://doi.org/10.1016/j.concog.2010.11.010

Markov, Y. A., Tiurina, N. A., & Utochkin, I. S.Different features are stored independently in visual working memory but mediated by object-based representations. *Acta Psychologica*, *197*, 52–63. https://doi.org/10.1016/j.actpsy.2019.05.003

Markov, Y.A., Utochkin, I.S., & Brady, T.F. (2019). manuscript). Real-world objects are not stored in holistic representations in visual working memory. Accessed on *PsyArXiv*, 20 September, 2019 version. https://doi.org/10.31234/osf.io/rkpnq

McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.

Mitchell, J. F., Sundberg, K. A., & Reynolds, J. H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron*, *55*, 131–141. https://doi.org/10.1016/j.neuron.2007.06.018

Mody, S., & Carey, S. (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition*, *154*, 40–48. https://doi.org/10.1016/j.cognition.2016.05.012

Moore, C. M., Stephens, T., & Hein, E. (2020). Object correspondence: Using perceived causality to infer how the visual system knows what went where. *Attention, Perception, & Psychophysics*, *82*(1), 181–192. https://doi.org/10.3758/s13414-019-01763-y

Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology*, *8*, 317–336. https://doi.org/10.1007/s13164-016-0311-6

Odic, D., Libertus, M. E., Feigenson, L., & Halberda, J. (2013). Developmental change in the acuity of approximate number and area representations. *Developmental Psychology*, *49*(6), 1103–1112. https://doi.org/10.1037/a0029472

Papenmeier, F., Meyerhoff, H. S., Jahn, G., & Huff, M. (2014). Tracking by location and features: Object correspondence across spatiotemporal discontinuities during multiple object tracking. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(1), 159–171. https://doi.org/10.1037/a0033117

Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: MIT Press.

Peacocke, C. (2001). Does perception have a nonconceptual content? *The Journal of Philosophy*, *98*(5), 239–264.

Pepperberg, I. M., Gray, S. L., Cornero, F. M., Mody, S., & Carey, S. (2019). Logical reasoning by a Grey parrot (*Psittacus erithacus*)? A case study of the disjunctive syllogism. *Behaviour*, *156*, 409–445. https://doi.org/10.1163/1568539x-00003528

Pollatsek, A., Rayner, K., & Collins, W. E. (1984). Integrating pictorial information across eye movements. *Journal of Experimental Psychology: General*, *113*(3), 426–442. https://doi.org/10.1037/0096-3445.113.3.426

Pollatsek, A., Rayner, K., & Henderson, J. M. (1990). Role of spatial location in integration of pictorial information across saccades. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(1), 199–210. https://doi.org/10.1037/0096-1523.16.1.199

Porot, N. (2019). Some non-human languages of thought. PhD dissertation, The Graduate Center, CUNY.

Prinz, J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.

Prinz, J., & Clark, A.Putting concepts to work: Some thoughts for the twenty-first century. *Mind & Language*, *19*(1), 57–69. https://doi.org/10.1111/j.1468-0017.2004.00247.x

Pylyshyn, Z. (2004 1984). *Computation and cognition*. Cambridge, MA: MIT Press.

Pylyshyn, Z. (2003). *Seeing and visualizing: It's not what you think*. Cambridge, MA: MIT Press.

Quilty-Dunn, J. (2019a). Is iconic memory iconic? *Philosophy and Phenomenological Research*, https://doi.org/10.1111/phpr.12625

Quilty-Dunn, J. (2019b). Perceptual pluralism. *Noûs*, https://doi.org/10.1111/nous.12285

Quilty-Dunn, J., & Green, E. J. (Unpublished). Perceptual attribution and perceptual reference

Quilty-Dunn, J., & Mandelbaum, E. (2018). Inferential transitions. *Australasian Journal of Philosophy*, *96*(3), 532–547. https://doi.org/10.1080/00048402.2017.1358754

Rescorla, M. (2009a). Cognitive maps and the language of thought. *British Journal for the Philosophy of Science*, *60*(2), 377–407. https://doi.org/10.1093/bjps/axp012

Rescorla, M. (2009b). Chrysippus' dog as a case study in non-linguistic cognition. In R. W. Lurz (Ed.), *The Philosophy of Animal Minds* (pp. 52–71). Cambridge: Cambridge University Press.

Rescorla, M. (2009c). Predication and cartographic representation. *Synthese*, *169*, 175–200. https://doi.org/10.1007/s11229-008-9343-5

Rugani, R., Regolin, L., & Vallortigara, G. (2008). Discrimination of small numerosities in young chicks. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*(3), 388–399. https://doi.org/10.1037/0097-7403.34.3.388

Scholl, B. J., & Leslie, A. M. (1999). Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In E. Lepore & Z. Pylyshyn (Eds.), *What is cognitive science* (pp. 26–73). ?Oxford: Blackwell.

Scholl, B. J., Pylyshyn, Z. W., & Franconeri, S. (1999). When are spatiotemporal and featural properties encoded as a result of attentional allocation? *Investigative Ophthalmology and Visual Science*, *40*(4), S797.

Sellars, W. (1956). Empiricism and the Philosophy of Mind. *Minnesota studies in the philosophy of science*, *1*(19), 253–329.

Shepherd, J. (2018). Intelligent guidance of action and the use of mixed representational formats. *Synthese*, https://doi.org/10.1007/s11229-018-1892-7

Shepherd, J. (2019). Skilled action and the double life of intention. *Philosophy and Phenomenological Research*, *98*(2), 286–305. https://doi.org/10.1111/phpr.12433

Smortchkova, J., & Murez, M. (forthcoming). Representational kinds. In J. Smortchkova, T. Schlicht, & K. Dolega (Eds.), *Mental Representations: The foundations of cognitive science*. Oxford: OU.

Solvi, C., Al-Khudhairy, S. G., & Chittka, L. (2020). Bumble bees display cross-modal object recognition between visual and tactile senses. *Science*, *367*, 910–912. https://doi.org/10.1126/science.aay8064

Stepper, M. Y., Moore, C. M., Rolke, B., & Hein, E. (2019). The role of object history in establishing object correspondence. *Attention, Perception, & Psychophysics*, *82*, 1038–1050. https://doi.org/10.3758/s13414-019-01923-0

Toribio, J. (2011). Compositionality, iconicity, and perceptual nonconceptualism. *Philosophical Psychology*, *24*(2), 177–193. https://doi.org/10.1080/09515089.2011.558836

Uller, C., Carey, S., Huntley-Fenner, G., & Klatt, L. (1999). What representations might underlie infant numerical knowledge? *Cognitive Development*, *14*, 1–36. https://doi.org/10.1016/s0885-2014(99)80016-1

Uller, C., & Lewis, J. (2009). Horses (*Equus caballus*) select the greater of two quantities in small numerical contrasts. *Animal Cognition*, *12*, 733–738. https://doi.org/10.1007/s10071-009-0225-0

Wang, B., Cao, X., Theeuwes, J., Olivers, C. N. L., & Wang, Z. (2017). Separate capacities for storing different features in visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(2), 226–236. https://doi.org/10.1037/xlm0000295

Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, *30*, 111–153. https://doi.org/10.1006/cogp.1996.0005

Xu, F., Carey, S., & Quint, N. (2004). The emergence of kind-based object individuation in infancy. *Cognitive Psychology*, *49*(2), 155–190. https://doi.org/10.1016/j.cogpsych.2004.01.001

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

How to cite this article: Quilty-Dunn J. Concepts and predication from perception to cognition. *Philosophical Issues*. 2020;1–20. https://doi.org/10.1111/phis.12185