Observational concepts

Daniel A. Weiskopf

1. From perception to thought

How is it that we are able to think about what we perceive? More specifically, how are we able to bring the resources of conceptualized thought to bear on the objects and events that are represented to us in perception? And how much of our capacity for conceptualized thought is undergirded by, or is an extension of, our capacities for perception and action? Addressing these questions requires disentangling some of the more tightly woven strands linking perception, thought, and action.

While one of the most distinctive things about human concepts is that they extend over indefinitely many types of things that transcend perception, we can also reflect on straightforwardly perceivable objects, and the concepts we have of these objects are often acquired by perceptually encountering and manipulating them. Additionally, how we perceive the world is infused or colored by the conceptual capacities that we possess. We don't merely see the cat, we see her *as* a cat, a visual state that depends on conceptualizing her in a certain way. And in virtue of seeing her as a cat, we may come to form perceptual beliefs concerning her presence and qualities. Thus, conceptualized perception enables conceptualized thought.

My aim here is to illuminate what happens at the interface between perception and higher cognition. On the view I develop, *observational concepts* are the pivot on which this relationship turns. Observational concepts are those that are spontaneously made available at the interface between perception-action systems and the conceptual system. They correspond to the ways we have of conceptually dividing the world based solely on the available perceptual information. We

are able to treat what is perceived as evidence for what isn't directly perceived (or, indeed, perceivable at all). Observational concepts form the evidential basis for these perception-transcendent inferences. And ultimately, we acquire ways of thinking about perceived things that are not tied directly to how they are perceived; observational concepts are central to this process insofar as they provide us with an initial way of conceptually tracking categories that we will learn to represent in perception-transcendent ways.

In what follows, I situate observational concepts in the larger architecture of cognition, characterize their role and content, describe how they are learned, and show how they play a key role in learning further concepts. Along the way I distinguish them from related constructs such as recognitional concepts and show that arguments against recognitional concepts fail to work against observational concepts. I conclude by discussing what observational concepts can teach us about the extent to which perceptual systems and representations may shape higher cognition.

2. Interfaces

Observational concepts are distinguished by their functional role, specifically by the location they occupy in the overall architecture of cognition. Many issues about cognitive architecture remain largely unsettled, but the only architectural assumption employed here is the distinction between input-output systems and central cognitive systems. Central systems include but need not be exhausted by the conceptual system, which is not assumed to be unitary.[1]

In the simplest case, the mapping from perceptual systems to central systems is direct, so that the hierarchy of levels of sensory processing culminates in contact between sensory systems

---

[1] The main views ruled out by this assumption are subsumption-style architectures and variations on them, i.e., any model on which there are mainly layers of horizontal connections running directly from sensory input to motor output, with few connections between these layers (Hurley, 2008). While I will mostly speak of *the* conceptual system here, massive modularists may replace this with talk about a host of central modules.

and systems of conceptual interpretation. There may, however, be complex perception-action

connections that bypass higher thought, as in some forms of online sensorimotor coordination.

On a more indirect arrangement, perceptual systems feed first into intermediate nonconceptual

systems that preprocess sensory inputs by transforming, modifying, tagging, and re-representing

them in various ways.[2] Whether the arrangement is direct or indirect, however, there must be

some sort of interface between the conceptual system and these various nonconceptual

processing systems. What goes on at such an interface is precisely a transition from the way the

world is represented in perception (or nonconceptually) to the way that we conceive of the

world.[3]

The outputs of perception, while nonconceptual, are enormously rich, presenting us with

a three-dimensional world that comes divided into distinct objects and events that are assigned

locations and motion vectors in space, along with various perceivable qualities: color, shape,

shading, texture, degree of solidity, etc. (Clark, 2004; Matthen, 2005). These outputs constitute

interface representations, inasmuch as they are both accessible to conceptual and nonconceptual

systems. The job of these representations is to present a limited body of information from one

system to another in a 'legible' format, where a representation's format is the way that its

semantic content is encoded in the structure of the vehicle. The way a system formats

information is tailored to the problem it solves or the task it carries out. Candidate formats for

---

[2] This is one way of viewing the systems of core cognition discussed by Carey (2009). Core cognitive systems are domain-specific mechanisms that analyze perceptual inputs and produce conceptual outputs, including foundational concepts of objects, number, and agency. While Carey holds that these systems have "integrated conceptual content" (2009, p. 67), she also holds that they use iconic but nonperceptual representations. Architecturally, they are an intermediate processing stage between perception and central cognition.

[3] I have been framing things here in a way that takes sides in the debate between conceptualists and nonconceptualists about perception. While there are many ways of drawing this distinction, for present purposes we can say that nonconceptualists hold that perceptual states are not individuated by reference to the conceptual repertoire of the cognizer, while conceptualists hold that they are. I am assuming nonconceptualism in the following sense: there is at least some range of perceptual states available to cognizers that is independent of the conceptual capacities that they possess; being in these states does not involve deploying any concepts.

mental representations include icons and images, propositions, frames, maps, models, analog accumulators, and so on.

Interfaces can be classified in several ways. First, systems can make use of the same representational vehicles or different ones. Consider a parsing mechanism that contains syntactic and semantic components. The syntactic parser might construct a tree that represents constituency, dominance, and other formal relations, which is then passed to the semantic parser for interpretation. Here the same representation is transferred from one system to another. Among interfaces that use the same vehicles, some involve passing along *all* of their representational outputs across the interface, while others involve passing only *some* of them. The latter case can arise where one system interfaces with two downstream systems, each of which has access to only some of its outputs. This occurs in vision, where after the common retinocortical pathway and processing in V1 the visual output divides into the dorsal and ventral streams, which make use of different properties of the visual input (Milner & Goodale, 2006). Where an interface involves passing all of one system's output to another, making use of the same vehicles, call this a *pass-through* interface; where only a subset of these representations are disseminated, call this a *filter*. Systems that compute intermediate level representations often impose filters that prevent these from being passed downstream, and some systems filter their outputs differently depending on the consumer systems that they are feeding.

Finally, some interfaces occur between systems that make use of different representational vehicles. Call these *translational* interfaces: ones that transform information encoded in one format to information encoded in another. The simplest example is a digitizer, which converts a continuous analog signal into discrete bits. Cognitive systems employ many different encoding schemes for different kinds of information, so there must exist at least some

translational interfaces of varying degrees of complexity. In numerical cognition there are mappings from a display of a quantity to an analog accumulator, in visually guided reaching a perception of an object in space is mapped to a motor representation of how it can be grasped, and in language production a semantic representation is mapped to a phonological word form. These systems involve different codes, so their interfaces are translational.

This way of thinking about interfaces has a nice virtue: it just happens to correspond well with the major historical proposals about how concepts are acquired from perception. Empiricist theories of concepts adhere to the *shared vehicles thesis*: conceptual thought re-uses the representational vehicles deployed in perception (Barsalou, 1999; Prinz, 2002). On classical empiricist views such as Locke and Hume's, the relationship between percepts and concepts is described in several ways, the simplest of which involves the process of copying. In Hume's terms, ideas are less vivid copies of impressions, where copying preserves an impression's form but drains some of its "force and vivacity". Concepts are copies of percepts that play a distinctive functional role in cognition, and since copying preserves most of the properties of representational vehicles, acquiring a concept from experience involves a pass-through mechanism that causally reproduces those vehicles in a different system.

Beyond simple Humean copying, Locke sometimes describes the acquisition of complex ideas as a kind of abstraction from experience with particular category members. A Lockean abstraction mechanism selectively deletes distinctive features and retains common ones, and in this sense it is a filter for these features. The output of this process is a general concept: a stripped-down perceptual representation that captures the characteristics shared by the instances and other category members. In much this way, Locke describes the learning of a general idea like HUMAN as a process of comparing ideas of individual persons (Mama, nurse, etc.) and

subtracting their distinctive qualities. As an approximation, then, we can think of Humean copying as a pass-through mechanism, and Lockean abstraction as a filter.[4]

Non-empiricist views of concepts, by contrast, see this interface as being translational, since they hold that thought employs an amodal code distinct from those used in perception. So Fodor (2008) distinguishes between perceptual and conceptual representations on the grounds that the former are *iconic* while the latter are *discursive*. In iconic representations, every part of the representation stands for a part of what is represented, while discursive representations have a 'canonical decomposition' into constituents, which are the only parts of the representation that contribute to its overall semantic interpretation. Photographs are iconic: in a photo of a koala, every part of the photo represents some part of the koala; the part containing its round, fuzzy ears represents those very ears, their shape and texture, etc. Sentences, on the other hand, are discursive, so that in the sentence "the koala has round and fuzzy ears", the string "the koala" represents a contextually specified bear but the string "has round and" represents nothing, since it is not a semantically relevant part. While every part of an icon is semantically significant, not every part of a discursive representation is. Accordingly, in moving from perceptual/iconic representations to conceptual/discursive ones, a semantically homogeneous input is discretized by a translation mechanism.[5]

The minimal conception of an interface is a device that mediates information transfer and control between two or more systems by letting the output states of the producer systems determine the input states of the consumer systems. An interface is a 'normal route' for this kind

---

[4] This is a simplification in many ways, not least of which that Locke seems to have had several theories of conceptualization running at once. See Gauker (2011, Ch. 1) for an excellent discussion.
[5] One needn't hold that concepts are encoded discursively to think of this interface as translational. Mental model-style views of higher cognition might employ maplike representations that nevertheless differ from the iconic representations employed in perception. See also Mandler (2004).

of configuration. And the taxonomy of possible interfaces is one that maps onto the historically

dominant accounts of conceptualization in the empiricist and rationalist traditions.

3. Observational concepts introduced

We can now reintroduce the idea of an observational concept. Where perceptual systems

make contact with central cognition, there is an interface whose function is to generate a

conceptual output in response to that type of perceptual input. Given the structure of the interface

and the right background conditions, certain percepts will be causally sufficient to produce

conceptual output states in a way that is not mediated by any other concepts.[6] Observational

concepts are those concepts that are spontaneously activated or made available for use solely on

the basis of a subject's being in a certain perceptual state.

Such concepts are the basic constituents of perceptual judgments such as the belief that

there is a cat on the green chair by the window, that those are sunflowers, or that there was just a

loud bang to my left. These are our perceptually informed opinions concerning the objects and

events that surround us; they are the judgments that we can be in a position to make about the

environment just on the basis of what perception plus our repertoire of observational concepts

make available.[7]

There are two counterfactual aspects to this specification of observational concepts. First,

these concepts are not *invariably* tokened when the right perceptual inputs are present.

Perceiving a tree does not always lead to tokening TREE, still less to thinking THAT IS A TREE. The

---

[6] Any such mediation, should it exist, cannot go on forever—if one concept's activation is itself conceptually mediated, then that mediating concept is observational. Given the nature of interfaces, we can know that this regress must terminate, since there cannot be infinitely many representations lying along this causal pathway.

[7] For an excellent discussion of the epistemology of perceptually basic beliefs in this sense, see Lyons (2008). The view being developed here does not purport to have any special epistemological status or import, although it might be developed in those directions.

reference to background conditions becomes important. These conditions have to do with resources like attention and memory as well as motivational factors, goals, interests, and other ongoing thoughts one is entertaining. Perceptual inputs are sufficient for activating the appropriate observational concepts so long as the background conditions are those that dispose the creature to be conceptually responsive to its perceptual input. This is compatible with the possibility that perceptual states only ever actually lead to thoughts involving such concepts against a background of already active thoughts, plans, and broader aspects of one's mental 'set'. It only needs to be *possible* for these concepts to be entertained in isolation from this surrounding mental activity.

Perceptions don't dictate what perceptual judgments follow from them, nor how the perceived scene will ultimately be conceptualized.. A trivial example: the visual perception of the cat dozing on the chair may give rise to the belief that the cat is on top of the chair or the (distinct) belief that the chair is under the cat. The cat herself might be represented as a cat, as an animal, or as Sophie. The same auditory perception might be equally well conceptualized as a sigh or a snort. All that is required is that some concepts are immediately made available by the occurrence of a perceptual state, though a representation's being made available, in the present sense, doesn't entail its being deployed for any particular purpose.

Second, perceptual input itself is not *necessary* for the activation of an observational concept, since being a concept requires availability for many other cognitive processes, including offline reasoning and planning, direction of other actions such as the formation of mental images, and so on. These tasks can take place under endogenous control, in the absence of any ongoing perception. So observational concepts have a distinctive set of application conditions that tie

them to perception, but they also enjoy the free deployability characteristic of concepts generally.

Observational concepts are not merely *occasioned* by perceptual experience but are in a stronger sense *directed at* or *applied to* what is perceived. In a weak sense, a percept can be associatively linked with any sort of mental or behavioral state. There are direct behavioral associations such as the photic sneeze reflex, or potentially arbitrary psychological associations such as thoughts of springtime coming to mind whenever I see lilies. In order to forge a link stronger than mere association between perceiving an object and formation of perceptual judgments about it, there needs to be a semantic relationship of co-application or co-indexing. A percept and concept are *co-indexed* when the concept is, or is disposed to be, applied to the very same object that the percept is tracking.[8] A function of interfaces is to provide such semantic relationships so that information about the same object can be collected across different cognitive systems and representational formats. When vision or any other object-centered sense divides the perceived environment into an array of entities, they are passed on to the conceptual system along with individuating tags that distinguish them as separate possible objects of classification and predication.

Finally, this division of representational labor gives us an account of recognitional seeing, or "seeing-as". To see a tree, and to see it *as* a tree, are two different states. The latter is more complex, in that it contains the former. Seeing x involves having a representation of the perceptual qualities of x that is indexed to the object x itself; seeing x as F involves having this

---

[8] Much work in mid-level vision has emphasized the need for such co-indexing representations (Pylyshyn, 2001; Scholl, 2001). This work is obviously relevant here, with the caveat that the instantiation tags referred to in this literature are usually opened when objects move, change, or otherwise become salient. I would expand this notion so that any object that can be distinguished from its environment can be tagged and be the subject of conceptual predication. In fact, co-indexing relationships are needed for two cognitive purposes. One is to ensure that multiple representations track the same object across processing systems and contexts. The other, emphasized by Clark (2004), is to ensure that separate representations are linked synchronically into a single representation of an object; i.e., to solve the "binding problem".

percept co-indexed with the concept F. States of seeing-as are thus perceptual-conceptual hybrids, having one foot in each system. But notice that seeing x as F falls slightly short of seeing that x is F. Seeing-as is a (partially) conceptualized state, but not a propositional state. Activating a concept, making it ready for use, is not the same as employing it in a judgment, which is the form of seeing-that. In this sense, you can recognize x to be F without entertaining the proposition that x is F.

4. Against perceptual inflation

Observational concepts are limited by the perceptual similarities we can reliably track. In particular, it would seem we cannot have observational concepts of abstract entities (those that have no physical, perceivable manifestations at all) and categories that have overly heterogeneous perceptual presentations. Call this, following Jesse Prinz (2005), the *Imperceptibility Thesis*. Prinz denies the thesis, and argues that we can perceive indefinitely many abstract categories; indeed, virtually anything we can form a concept of can be perceived. This includes abstract categories of things such as numbers, or properties such as *being a philosophy major* or *being an uncle*.

Prinz's argument turns on his view about recognitional perceiving. One perceives X just in case: (1) X impinges on one's sensory transducers, (2) leading to the formation of a perceptual representation, (3) which is in turn matched against stored representations of X. So to perceive something involves forming a percept of it and retrieving a representation of it, where the content of the retrieved representation determines the content of the perception one has. In addition, Prinz follows Dretske's (1986) informational semantics, on which M represents C just in case instances of C have the power to cause instances of M, and this relationship holds because M has

the function of detecting C. Content is a matter of teleologically established and sustained informational connections.

Many abstract and intuitively imperceptible qualities will turn out to be perceivable on this view, since there is always some perceptual scenario that can cause us to retrieve representations of them. Consider being an uncle. Uncles have no common perceptual look, so I cannot just pick them out in a crowded room. However, if I arrange for all and only the uncles to raise their hands (assuming everyone is sincere and cooperative), they suddenly share such a look. If I arrange to get myself to think someone is an uncle just in case their hand is raised, and hand-raising is a detectable look, then what I am perceiving is *unclehood*, since my UNCLE concept is activated when I see their raised hands.

Now consider number. Small numbers might be perceivable; triples of discrete objects have a distinctive look, so perhaps we can perceive *threeness* and form the corresponding observational concept. Few think that we can perceive or form such concepts of higher numbers such as fifty-seven. But we can arrange a situation in which we will activate FIFTY-SEVEN in response to a perceptual scenario: we simply count the objects, keeping track of them using number words. Our saying the numerals out loud or internally is a perceptual state, so when we run out of things to count we have perceived the number corresponding to the last numeral that was perceptually activated. Generalizing these examples shows that we are able to perceive both abstract entities and qualities, so the Imperceptibility Thesis would be false.

But Prinz's account of perceptual content is implausibly permissive. Neither perception nor observational concepts themselves have contents as rich as the arbitrarily abstract categories that we can conceive of. The problem lies with the open-ended notion of recognitional perceiving. Recall that *any* sort of representation that can be retrieved as part of the comparison

or retrieval process initiated by perception can contribute to what is perceived, and there are no limits on how complicated these processes may be. This lack of limits is what lets us perceive uncles by seeing raised hands—but since this connection depends on a complex cognitive arrangement that goes beyond what the perceptual system itself contributes, we should not consider it an act of perception, properly speaking. Similarly in the number case: the act of counting large numbers uses more cognitive resources (memory, language, etc.) than just what our perceptual systems provide.[9] Worries about this form of content inflation are blocked on the account of observational concepts given here by the requirement that their deployment be under the control of an interface, rather than allowing any downstream representation to count.

Susanna Siegel (2010) has also argued that the contents of perception are extremely rich, so rich that we can perceptually represent not just colors, shapes, textures, and the rest of the qualities that traditional theories of vision admit, but also so-called "K-properties". These include properties covering objects (natural and artifact kinds), actions and events, mental states, and semantic facts, for example: *person, mountain, bicycle, carrying a dog, being inquisitive, being a word of Russian, being a phrase meaning 'the highway exit is ahead'*. According to the Rich Content View, all of these are possible contents that we can grasp in experience, and we can recognize this fact from attending to the phenomenology of the relevant experiences.

Siegel's argument runs as follows. Consider two related experiences: in E1, you are looking at a bunch of pine trees, but you do not have the ability to visually recognize pine trees; and in E2, you are looking at the same trees, but you have learned to recognize them. There is, it

---

[9] Prinz does argue separately that our sensory systems proper can represent these abstract categories as well. Here he relies on the existence of downward projections from higher cortical centers to the sensory systems themselves. These downward projections allegedly allow conceptual information to "infuse" sensory processing. However, on Prinz's official theory of content it is hard to see how this might work. If the dominant cause of activity in sensory systems is the perceivable object, then that is what sensory representations refer to, not abstractions. If their cause is split among external objects and higher level systems, the theory does not assign them content, since their function and informational role is indeterminate.

would seem, a phenomenological difference between these two experiences. Being able to perceptually recognize what you are looking at is different than not knowing what you are looking at, and the overall phenomenological state one occupies in each case seems different. To arrive at the Rich Content View, three further premises are needed: (1) if E1 and E2 differ in phenomenology, then there is a phenomenological difference between the experiences themselves; (2) if there is a phenomenological difference between them, then they differ in content; and (3) if they differ in content, it is a difference with respect to the K-properties that those experiences represent. Since we have granted that E1 and E2 differ in phenomenology, we quickly reach the conclusion that experience represents K-properties.

The account of how concepts come into play in acts of perceptual recognition gives us a response to this argument, one that focuses on the denial of premise (1). We should grant that there is a phenomenological difference for the perceiver having E1 vs. E2, but deny that this difference is located in the visual experience proper.[10] Siegel, in sketching possible responses, considers two related objections to this premise. Both involve saying that the phenomenological difference between E1 and E2 lies not in the visual experience itself but in some associated psychological aspect, in particular in some sort of *cognitive* experience. If E1 and E2 differ only in this related respect, there is no difference in visual content, particularly not with respect to the representation of K-properties.

The candidate cognitive experiences Siegel considers, however, are all forms of propositional representation. They are either commitment-involving attitudes (e.g., beliefs or judgments), or else noncommittal attitudes (e.g., hunches or merely entertained thoughts).

---

[10] I should add that talk of experiences here is Siegel's; I am glossing 'visual experiences' as being states that depend for their instantiation on activity in visual systems; similarly, 'cognitive experiences' are states that depend on non-perceptual systems in some way. This gives a way of mapping experience-talk onto talk of underlying cognitive systems.

Against the idea that the phenomenological difference comes from the presence of a commitment-involving attitude such as the belief that *that* is a pine tree, notice that I might well believe (because I am reliably informed) that the tree-appearances I perceive are not real trees, but only props or holograms. Despite this belief, they might appear phenomenologically different in situations E1 and E2. So committal attitudes such as beliefs cannot make the difference between the two.

Against the idea that noncommittal attitudes might explain the difference, she presumes these attitudes to be occurrent thoughts such as *that kind of tree is familiar*. Thoughts of seeming familiarity say nothing about what the objects perceived themselves are. However, she argues, these thoughts are simply unnecessary: "[t]here need not be, it seems, an extra episode (or occurrent state), beyond sensing, for the phenomenological change to take effect." (p. 106). So any alleged noncommittal attitude would be redundant in explaining the E1/E2 difference. Given that (with a few caveats) these two possibilities exhaust the primary forms of non-sensory experience, we can conclude that the phenomenological differences are located in the visual experiences themselves, not in any adjoining cognitive states.

Observational concepts have a distinctive functional role that renders them well-poised to thread this needle, however. First, as noted above, they are not propositional representations, and they do not *inherently* carry committal force. They may do so, if a concept becomes activated to the point where it is actually endorsed and applied to a perceived scene. The application of a concept to perception is a state that has distinctive correctness conditions, and thinkers are committed to these. But since activation is graded, these concepts may be exerting an effect even though they are not past the threshold for being applied. In the hologram case, once I am told that

the trees are fake, I refrain from applying the concept, but it may retain a residual level of activation.

What, then, of the objection that these concepts constitute merely a redundant noncommittal component? It is not clear that we have a separate occurrent state here at all. Observational concepts are distinct from their perceptual antecedents, true, but their activation levels are gradual. So the simple notion of a state's being (non-)occurrent does not obviously have application here. Certainly the issue cannot be settled by intuitions about how many "episodes" are involved in an experience. Siegel suggests that if a non-sensory event is not explicit and occurrent, then "it becomes less clear that it is phenomenally conscious at all" (p. 107). But why is this? The phenomenal properties of cognitive states, or their contribution to overall phenomenology, may in principle depend on any aspect of their functional or representational role. To take an example, tip-of-the-tongue states are cognitive, but they have a highly distinctive phenomenology that is plausibly underpinned not by explicitly entertaining any proposition, but rather by cascades of activation washing through networks of lexical and semantic memory. If these types of spreading activation can contribute to cognitive phenomenology, the same should be true of the varying levels of activation in the case of observational concepts.

What these two cases show is that attending to the role of observational concepts can help to establish boundaries on the representational power of perception. Architectural divisions determine what can be recognitionally perceived, thus blocking the sort of content inflation Prinz's account invites. Similarly, a wider range of functional interactions between perception and cognition can help us to account for the phenomenology of perceptual recognition without

endorsing Siegel's Rich Content View.[11] The phenomenological evidence that Siegel draws our attention to needs to be explained, but only by uncovering the details of the larger architecture and how it assigns content to underlying systems and states.

5. Observational vs. recognitional concepts

Observational concepts in my sense are not the same as recognitional concepts. Fodor presents a view of recognitional concepts according to which:

a concept is recognitional iff: (1) it is at least partially constituted by its possession conditions, and (2) among its possession conditions is the ability to recognize at least some things that fall under the concept as things that fall under the concept. (1998, p. 1)

Recognitional concepts are those that are in part constituted by abilities to recognize some things in their extension. If RED and SQUARE were recognitional, then possessing them would mean necessarily being able to apply them to red things and squares, respectively, on the basis of the appropriate perceptual encounter.

Fodor argues that there cannot be any such recognitional concepts, on the grounds that recognitional abilities fail to be compositional. Since compositionality is a necessary condition on concepts, anything non-compositional cannot be part of what individuates concepts. The main premise of his argument states that the possession conditions for a complex concept must include the possession conditions for its constituents. So whatever states, capacities, dispositions, etc. are required for having complex concepts such as RED SQUARE, FISH FOOD, or ONE-ARMED MAN necessarily include whatever is required to have the concepts that make them up. Possession conditions are inherited compositionally. But the ability to recognize instances is not inherited in

---

[11] Although I am not necessarily endorsing the idea that all of Siegel's K-properties can be the content of observational concepts, in my sense. Addressing this issue would require more clarity both on the scope of observational concepts, and on how to fill in the list of K-properties.

this way. One can be able to separately recognize fish and food but not be able to recognize that those confetti-like flakes are fish food. Recognizing fish food depends on knowing something about what fish actually eat, not on anything one could extract from the constituent concepts themselves. Since anything that is a constitutive property of a concept must be inherited compositionally by its hosts, recognitional abilities cannot be concept-constituting.

This argument doesn't work, however.[12] Suppose that some concepts have recognitional capacities c_1-c_n among their essential properties, and that these capacities must be among the possession conditions for any complex concept that hosts them. Suppose too that there is no corresponding recognitional capacity c^* for the complex concept itself—or if there is one, it is neither among c_1-c_n, nor is it derivable from them alone. This shows that not every concept composed from recognitional concepts is itself thereby recognitional. Even so, recognitional capacities are 'inherited' in the relevant sense—having a complex concept entails having the capacities that are the possession conditions for its constituents, and this helps to explain why having the constituents is roughly sufficient for having the complex concept. So the fact that a wholly distinct recognitional capacity for fish food fails to be produced compositionally from the combination of FISH with FOOD doesn't argue against the existence of recognitional concepts, so long as having FISH FOOD entails having the appropriate recognitional capacities for fish and food, taken individually.[13]

---

[12] The argument given here is similar to one advanced by Recanati (2002). For careful analysis and criticism of the notion of a recognitional concept, see Dahlgrun (2011); for a Fodorian defense of one type of recognitional concept, see Rives (2010).

[13] Fodor does in fact think that "people who have the concept PET FISH do generally have a corresponding recognitional capacity" (1998, p. 8). It just can't be one they've derived solely from the recognitional capacities of that concept's constituents. But this is only a problem if one wants to maintain that all complex concepts formed from recognitional concepts are also recognitional, which there is no reason to do. Fodor argues that this hybrid view is arbitrary or theoretically inelegant, but the architectural considerations pointed to above show that it is, in fact, quite predictable.

The model of observational concepts gives a principled explanation for why we would not expect recognitional abilities to compose in Fodor's sense. An observational concept for FISH allows one to categorize things as fish, as long as there are good instances presented in good perceptual conditions. However, this ability is mediated by the structure of the interface, which takes a restricted set of percepts as input. Imagine the interface as consisting of a set of perceptual analyzers attuned to category appearances. These perceptual analyzers may respond to what counts as good instances of fish, and so similarly would analyzers that take appearances of food as input. But the existence of these two analyzers entails nothing about the existence of a *third* device for responding to the characteristic appearance of fish food; in fact, given that it looks rather un-foodlike, one would predict that there isn't any such device, and hence that there is no observational concept FISH FOOD.

In any event, however, observational concepts are unlike Fodorian recognitional concepts, and also unlike Peacocke's (1992, Ch. 3; 2001) perceptual concepts in that their identity *as concepts* is not constituted by their perceptual connections. Observational concepts have a certain primary perceptual affiliation, namely a means of being directly deployed by perceptual inputs. This affiliation makes them observational, but it is not part of their possession conditions qua concepts.

Rives (2010) has defended a Fodorian notion of recognitional concepts. He argues that causal links to perception are essential to some concepts, on the basis of the general principle that scientific kinds are individuated by the laws that they participate in. For concepts, these laws include those that fix their content and causal role by connecting them with particular perceptual states. So principles of kind individuation mandate that some concepts are recognitional. Undoubtedly the general point that kinds in the special sciences are taxonomized by their causal

and functional profile is correct. But many of these causal links are *fragile*, and we should be wary of making concept identity contingent on them.

For an actual case in which observational concepts have their perceptual affiliations severed, consider associative agnosia. The core disorder involves the inability to identify visually presented objects either by naming or by nonverbal means (e.g., sorting). Patients tend to perform best with real objects or highly realistic ones such as photographs; when they hazard an identification, they often confuse visually similar objects. Associative agnosics can sometimes describe the features that objects have, but seem not to know *what* the described objects are. The disorder is multifaceted, and its basis is not entirely clear (Farah, 2004, Ch. 5; Humphreys & Riddoch, 1987; 2006), but its existence indicates that specifically visual routes to conceptual deployment can be disrupted without corresponding loss of the concepts themselves.

The patient C.K., for example, was capable of providing elaborate verbal descriptions of objects that he could recognize by touch but not by sight. He could also answer specific questions about the visual properties of objects, suggesting that this fine-grained visual knowledge was usable in reasoning tasks but not in object identification (Behrmann, Moscovitch, & Winocur, 1994). The patient ELM, who played in a military brass band, was able to freely describe the nonvisual properties of brass instruments in detail, and to use concepts of these instruments in an associative learning task despite his deficits in visual identification (Dixon, Desmarais, Gojmerac, Schweizer, & Bub, 2002). The converse pattern also appears: in some cases of category-specific semantic impairments there is loss of almost all information that might be used in reasoning about a category, despite relatively spared ability to identify category members (De Renzi & Lucchelli, 1994). So the capacity to identify category members and the capacity to reason about them are at least partially separable.

These cases suggest that agnosic patients have not *lost* the concepts of the objects that they can no longer identify, and that concepts acquired observationally can persist once these core functional links are severed. Indeed, it seems to be a general truth that concepts are *not* particularly fragile; they can survive arbitrarily many gains and losses in their representational affiliations. As a design feature, this makes sense: the more such essential links a concept has, the more difficult it is to acquire and the easier it is to lose. Concept possession should be buffered from such changes in cognitive structure, however. Given this, the notion of an observational concept seems to be a marked improvement on both the notion of a recognitional concept (as Fodor conceives it) and of a perceptual concept (as Peacocke conceives it). Both of these notions commit us to constitutive links between concepts and perception, but these, I have argued, we have ample reason to reject.

6. Learning observational concepts

Observational concepts link the conceptual system to the world in a flexible, open-ended way. Learning an observational concept involves two stages: first, the construction of a dedicated perceptual analyzer that can co-classify objects based on their appearances; second, the construction of a link between the output of that analyzer and a symbol in the conceptual system.

The main job of a perceptual analyzer is to create structure in a perceptual similarity space. There are many mechanisms for achieving this. Local, bottom-up processing has been most intensely studied in theories of object recognition, where the debate has focused on whether multiple stored views of an object or single viewpoint-independent structural descriptions are required (Feldman, 2003; Riesenhuber & Poggio, 2000; Pessig & Tarr, 2007; Tarr, 2003). These representations also need to be processed by the right sort of generalization mechanism or

similarity gradient, along with a weighting matrix for assigning parts of the object greater or lesser significance. Once a set of views of an object have been stored and linked together, or once a structural description has been generated, the process of analyzing new percepts using this device depends just on the fit between the percept and the stored representations.

However, perceptual analyzers may be constructed from many different materials. How an object fits into a larger perceived scene is also relevant. Objects that appear in a familiar location, or in the context of other familiar objects, tend to be recognized faster and more accurately (Henderson & Hollingworth, 1999). Analyzers take into account holistic scene-level properties in determining object categorization, not just local qualities. More interestingly, object recognition also depends on whether something is appropriately poised for action (Humphreys & Riddoch, 2007). Recognition can be facilitated by the orientation of an object relative to the normal actions one performs with it, and its orientation relative to other objects in a perceived scene, even if they are not the usual ones that it is seen to interact with.[14]

Perceptual analyzers, then, may have extremely rapid access to a large amount of information. They can integrate stored viewpoints and structural descriptions, as well as holistic scene-level properties and associations and stored or potential patterns of motor interaction and event schemas that objects can enter into. Neurally, the implementation of visual analysis seems to involve such widespread retrieval insofar as it taps massively sweeping feedforward projections as well as recurrent projections that play a role in mediating visual awareness (Lamme, 2006; Lamme & Roelfsema, 2000). These analyzers may be assembled relatively quickly—seven hours of training is sufficient to induce new category-specific responses in the

---

[14] I should add that the view of perception I am adopting here is compatible with the existence of many tight links between perception and action systems. Indeed, there is robust evidence that the two are tightly coupled, so that sensorimotor integration rather than separation might be the norm (Prinz, 1997; Schütz-Bosbach & Prinz, 2007). Observational concepts, nevertheless, occur where perceptual systems interface with concepts, no matter how richly connected perception and action themselves may be.

fusiform gyrus (Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999). None of this processing, however, needs to be conceptually mediated. Extraction of viewpoints and structural descriptions can occur prior to knowing what the object being analyzed is, as can scene analysis and co-occurrence detection. Action-based understanding may stem from keeping track of our own past actions with perceived objects as well as from on-line computation of the plausible acts and events that objects afford.

The traditional view about the scope of observational concepts that they are situated at the "basic level" (Jacob & Jeannerod, 2003, pp. 140-1; Rosch et al., 1976). Basic categories of objects such as cats, birds, hammers, airplanes, and squares can be grouped together on the basis of the similarities and differences detectable by the perceptual system itself, even at 3-4 months old (Eimas & Quinn, 1994; Quinn, Eimas, & Rosenkrantz, 1993). But narrower object categories also involve such similarities: ball-peen hammers versus claw hammers, MD-80s versus Boeing 747s, domestic shorthairs versus Manx cats. Expert object classification treats such specific groupings as entry points. Most of us behave like experts when it comes to classifying well-known individuals, as shown by the fact that we can rapidly and correctly identify Bill Clinton, Golda Meir, the Mona Lisa, or the Sears Tower across different encounters (Bukach, Gauthier, & Tarr, 2006; Tanaka, 2001; Tanaka & Gauthier, 1998). Even relatively global perceptual similarities may be detected early, such as those that unite all entities that move in the same way as animate living beings, or those that move in a way distinctive of vehicles (Behl-Chadha, 1996; Mandler & McDonough, 1993; 1998). Observational concepts naturally coalesce around such islands of perceptual similarity.

Having a set of perceptual analyzers causes perceived objects to cluster in similarity space, but it does not give one the ability to *think* about those objects until they are assigned to

conceptual representations (Gauker, 2011, Ch. 3). This requires noticing the similarities that are being detected by a perceptual analyzer and coining a new mental symbol for the purpose of keeping track of *that kind of thing*, namely the kind of thing that has the relevant appearance which the underlying analyzer tracks.

Conceptualization thus involves both a functional transformation and a shift in content. Perception knows about looks: it informs us about how things appear, and it tracks and produces detectable similarities among appearances. Concepts, in the first instance, are not for tracking appearances, but for determining what things are, and for representing them in a way that is autonomous from perception. So the content of perception informs us about the appearances of objects around us, but when we coin observational concepts they must be (as the agnosia cases show) potentially detachable from these appearances: they track categories of things that have, but are not defined by, those appearances. Perception groups things that *look F-like*, while concepts track things that *are Fs*. Part of what is involved in generating observational concepts from perceptual analyzers is shifting from content that represents only looks to content that represents categories as distinct from their looks.[15]

This point can be illustrated using Peacocke's account of perceptual content (1989, 1992). For Peacocke, the accuracy conditions for perceptual states are given in terms of abstract scenario types. A scenario is a way of filling out space in the vicinity of the perceiver; it specifies a way of locating surfaces and their properties relative to an origin point and a set of orienting axes. A perceptual scene, then, is a way that this spatial volume may be filled in. A filled-in

---

[15] One may object here that some perceptual representations seem to track more than mere appearances. Some models of object vision propose that there is a stage at which object-centered visual representations such as 3D models of object parts and structure are computed. These are independent of the particular viewpoint of an observer and hence may seem to be about more than looks. The difference, however, is that even these 3D models do not distinguish between being a member of one category that has a certain 3D profile and being a member of a different category that just happens to share that profile. That is, these 3D models are not yet *kind* representations, and hence are still about appearances in the generalized sense meant here. Thanks to Jake Beck for raising this objection.

scene depicts the world as containing surfaces at various distances and orientations, instantiating properties such as color, texture, shading, and solidity. This captures the look of the scene, but it does not specify what kinds of objects the world itself contains. The very same perceptual scene is compatible with many different conceptualizations. This is true not just in the sense that there may be various different entry-points into scene categorization, but also in the sense that however things happen to look in a scene, they may not be what they appear on their surfaces to be. The cat may be a puppet, the tree a prop. Our recognition of this "face-value gap" is possible because the function of concepts is to provide possibilities for representation and control of thought and behavior that go beyond what perception delivers.

Conceptual content, then, reflects the appearance-transcending role of the conceptual system. This point has nothing to do with the standard way of drawing the distinction between perceptual and nonconceptual states that relies on appeals to the fine-grainedness or richness of perceptual representation (Tye, 2005). There is no inherent reason that rich, fine-grained properties cannot enter into the content of some of our concepts. This is suggested by the fact that observational concepts present their content under a perceptual mode of presentation. Modes of presentation generally are ways in which categories are grasped in thought; they are *how* we are thinking about a particular category. An observational concept presents its content as being *the kind of thing that looks like THAT*, where *THAT* refers to the complex representational content that is computed by the relevant perceptual analyzer. Hence having the concept normally involves having access to facts about looks. Given this access, we are also capable of forming the concept of *THAT* sort of look.[16] We can either think about the things that the appearance is

---

[16] This point is also made by Peacocke (2001, p. 258): "When some property is given to a thinker in perception in way W, then, if that thinker is sufficiently reflective, there seems to be a conceptual way C_W of thinking of that way W, where this conceptual way C_W seems to be made available to the thinker precisely by her enjoying an experience whose nonconceptual content includes that way W."

presenting us with, or else we can think about a particular kind of appearance, the difference

being in the former case the reference of our thought 'reaches through' the appearances to what

possesses them, while in the latter case we are thinking about an appearance-defined category.

While observational concepts aim to track categories that have a certain appearance but also may

have an appearance-transcendent nature, the interface that enables them to do this job also makes

available concepts of appearances themselves.

The process of learning an observational concept bears little resemblance to traditional

hypothesis testing models of learning. For one thing, what is learned is not a hypothesis or a

judgment, but rather a two-part cognitive structure composed of a perceptual similarity detector

coupled to a symbol that it is capable of activating, but which is also free to function

autonomously in acts of higher thought. The learning process itself may involve largely passive

extraction of information from the environment, as when we store specific viewpoints on an

object or when we make associations between perception and action. Sometimes it may involve

task-relevant but incidental perceptual learning (Schyns, 1998; Schyns, Goldstone, & Thibault,

1998). Or, as in supervised learning studies, it may involve conscious, explicit learning of

categorical distinctions, accompanied by environmental feedback.[17]

Nevertheless, the processes involved in building analyzers are undeniably varieties of

learning—or rather, they are if we put aside our *a priori* prejudices about what learning must be

and focus on determining what learning as an empirical phenomenon actually is (Laurence &

Margolis, 2011). The process of producing these new cognitive structures is environmentally

sensitive in such a way that the product's detailed functioning depends on its history of

interactions. The representational properties of the structures themselves reflect the kinds of

---

[17] For discussions of the variety of category learning tasks and the neural bases involved in each of them, see Ashby & Maddox (2011); Segar & Miller (2010).

inputs that shaped them. And the structures acquired tend to be both long-lasting and adaptive, in the sense that they improve the organism's abilities to carry out practical and cognitive tasks. These are all stereotypical marks of learned psychological capacities (Landy & Goldstone, 2005; Goldstone, Gerganov, Landy, & Roberts, 2008).

7. Observational concepts and conceptual development

Once they are learned, observational concepts can resolve some longstanding puzzles in the theory of concepts.[18] One problem arises for theories that posit complex internal structure for concepts. The constituents of a concept are its *features*: they pick out the information about the concept's referent that the concept encodes. Prototypes, exemplars, causal models, and other complex types of representations are built up from such features and the relations among them. Most of the dominant models of conceptual structure depict them as being composed of features (Weiskopf, 2009b).

However, as Eleanor Rosch noted some time ago, this view may be developmentally implausible. Describing the features elicited in some of her studies, she says: "some attributes, such as 'seat' for the object 'chair', appeared to have names that showed them not to be meaningful prior to knowledge of the object as chair" (Rosch, 1978, p. 42). Moreover, "some attributes such as 'you eat on it' for the object 'table' were functional attributes that seemed to require knowledge about humans, their activities, and the real world in order to be understood." (p. 42). The problem, in short, is that many of the features displayed by adult concepts are ones that it is not easy to think of as being developmentally prior to the concepts that they are part of. Call this the *feature problem*.

---

[18] This general point is also endorsed by Roskies (2008), who argues that nonconceptual capacities are needed to explain concept acquisition. The account here aims to flesh out this picture, and also to show that concepts newly learned from perception can play further "bootstrapping" roles in development.

Conceptual atomists take the feature problem to be an objection to the view that concepts are internally complex. The fact that these features are clearly linked with concepts after they are already acquired shows that they cannot be *part* of those concepts, since coming to possess a complex concept requires first possessing its parts. Another possible interpretation, however, is that concepts are not developmentally fixed. The initial form that a concept takes is not necessarily the same as its mature form. Many concepts may start out life as simple observational concepts, later developing into complex, structured representations having many different constituents and inferential liaisons.

Suppose concepts undergo a kind of developmental cycle. They might originate as unstructured observational concepts attached to a complex perceptual analyzer. Having an observational concept allows one to attend to objects 'as such'—that is, to think of them as being unified groupings that are worth keeping track of, and which may have interesting properties and behaviors to discover. Observational concepts that track especially useful categories, and which are frequently deployed in important cognitive tasks, are then vehicles for directing our attention and intentional behavior towards a category. The more we attend to a category, the more things we discover about its members, including about their perceptual properties. So we may come to analyze the appearance of a certain type of animal into its parts (legs, tail, teeth), or its observable behaviors into classes (purring, grooming). As we notice these various perceivable qualities we can elaborate on our original observational concept and the perceptual analyzer that underlies it. Where it might have started life as an unstructured symbol, now it incorporates features corresponding to the various bits of information we have collected about the typical look of the animals it responds to. The fact that concepts undergo these types of changes in their structure and content is no surprise if we are already committed to the pluralist idea that the

conceptual system makes use of a variety of different types of representations in different domains, tasks, and contexts (Weiskopf, 2009a).

One aspect of conceptual development, then, might be repeated cycles of representational redescription or re-encoding such as this.[19] Concepts of types of objects might come first. An early TABLE or CAT concept might not include conceptualized information about how those entities are put together or how they function, but this information can come to be part of these concepts with time. Infant categorization, for instance, often seems to be based on perceptual sensitivity to particular parts and their structural configuration, but there is no reason to think they conceptualize these parts until later (Rakison & Butterworth, 1998a; 1998b). From object concepts emerge concepts of object parts, properties, and relations, so the fact that, for adults, TABLE has LEG as a feature says nothing about developmental priority. This kind of complex developmental pattern is consistent with the emergence of more complex concepts out of simple ones, and simple ones (ultimately) out of perception.

8. Conclusions

There is obviously much more to be done the flesh out the sketch of observational concepts provided here. I have spoken almost entirely about perception, and said nothing about concepts that interface directly with action systems, but these are important in developing a psychologically realistic account of basic actions. I have also not worked out in detail the precise developmental trajectory of observational concepts, but doing so is needed for understanding the

---

[19] Karmiloff-Smith (1992) also uses the term "representational redescription", but my usage differs from hers. I mean only that the very same content can be encoded in many functionally or formally distinct representations at different stages of learning or development. Thus we might start with a purely observational concept of a certain category and then later acquire a number of exemplars from it, a prototype for it, and so on for other complex cognitive structures. These are ways of redescribing or repackaging the same conceptual content. Karmiloff-Smith, by contrast, means something slightly more specific, namely that content initially only grasped implicitly will in time pass through a number of encoding stages to become both available to explicit conscious access and linguistic formulation.

nature of the so-called "perceptual-to-conceptual shift" in infancy (Rakison, 2005). Nor have I given an account of how language both adjusts the boundaries of perceptual similarity spaces and cues the creation of new concepts.

Despite these lacunae, observational concepts seem empirically well-grounded, theoretically plausible, and indeed, arguably necessary in understanding the origins of thought. They also suggest a larger moral. Recent debates over empiricism, and over the extent to which thought is grounded in perception more generally, have mostly focused on sweeping claims about the nature of all of our concepts. Thoroughgoing concept empiricism is unlikely to be true (Dove, 2009; Machery, 2007; Mahon & Caramazza, 2008; Weiskopf, 2007); but empiricism might be *almost* true for our observational concepts. I say *almost* here because observational concepts are not simple copies of percepts. As the feature problem indicates, not all of the structure in percepts makes it into our initial observational concepts. While this structure may be added with time, it requires reanalysis and the construction of new representations—something that is less characteristic of a Lockean filter, and more like a translation mechanism. Observational concepts are not quite copies of percepts, but they are the concepts that are shaped most powerfully by perception, and they are where the design of our perceptual systems leaves its imprint most visibly on thought.

Claims about thought being shaped or molded by perception fall far short of committing us to strong empiricist conclusions. On the other side of the coin, amodal theorists also need to acknowledge that these shaping effects may, for all anyone knows, be profound and far-reaching. If our first thoughts involve largely observational concepts, then this places us in a certain starting point in the space of possible developmental trajectories for conceptual systems to follow. Where we can go from that starting point, and what regions of conceptual space are

inaccessible to us, depends on the richness of the processes available to us for amplifying and refining our conceptual repertoire. Knowing the range of observational concepts we can entertain will help to locate human thought in the space of possible conceptual schemes.

This begins to give us a sharper picture of the research directions that open up once we adopt the notion of an observational concept. There has been an enormous amount of developmental work focused on concept acquisition and the many mechanisms that are implicated in it, but little consensus on how to describe the particular functional and representational changes involved in acquiring concepts from perception. We are now in a position to start fleshing out our sketchy understanding of how the conceptual system is integrated with other faculties in our cognitive architecture. Theories of concepts have so far focused to a large extent on internal affairs: what kinds of representations concepts are, how they are processed and deployed, and so on. These accounts are important, and they tell us much about the organization of the conceptual system *per se*, but little about how it interfaces with the host of other cognitive systems we possess, or how concepts as a distinctive type of representation help to orchestrate cascades of cognitive processing. Describing the nature of these interfaces is an open research problem, and one worth pursuing insofar as it promises to inform us about the large-scale cognitive architecture of which our distinctively human conceptual capacities form only a part.[20]

Acknowledgements

---

[20] For some suggestions on how to distinguish higher and lower cognitive faculties, and how conceptualized thought itself might be functionally defined, see Weiskopf (forthcoming).

References

Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annals of the New York Academy of Sciences*, 1224, 147-161.

Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-660.

Behl-Chadha, G. (1996). Basic-level and superordinate-like categorical representations in early infancy. *Cognition*, 60, 105-141.

Behrmann, M., Moscovitch, M. and Winocur, G. (1994). Intact mental imagery and impaired visual perception: Dissociable processes in a patient with visual agnosia. *Journal of Experimental: Psychology Human Perception and Performance*, 20, 1068-1087.

Bukach, C. M., Gauthier, I., & Tarr, M. J. (2006). Beyond faces and modularity: The power of an expertise framework. *Trends in Cognitive Science*, 10, 159-166.

Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.

Clark, A. (2004). Feature-placing and proto-objects. *Philosophical Psychology*, 17, 443-469.

Dahlgrun, M. (2011). The notion of a recognitional concept and other confusions. *Philosophical Studies*, 150, 139-160.

De Renzi, E., & Lucchelli, F. (1994). Are semantic systems separately represented in the brain? The case of living category impairment. *Cortex*, 30, 3-25.

Dixon M.J., Desmarais G., Gojmerac C., Schweizer T.A., & Bub D. (2002). The role of premorbid expertise on object identification in category-specific visual agnosia. *Cognitive Neuropsychology*, 19, 401-419.

Dove, G. (2009). Beyond perceptual symbols: A call for representational pluralism. *Cognition*, 110, 412-431.

Dretske, F. (1986). *Explaining Behavior*. Cambridge: MIT Press.

Eimas, P. D., & Quinn, P. C. (1994). Studies on the formation of perceptually based basic-level categories in young infants. *Cognitive Development*, 65, 903-917.

Farah, M. (2004). *Visual Agnosia* (2nd ed.). Cambridge: MIT Press.

Feldman, J. (2003). What is a visual object? *Trends in Cognitive Science*, 7, 252-256.

Fodor, J. (1998). There are no recognitional concepts, not even RED. *Philosophical Issues*, 9, 1-14.

Fodor, J. A. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Gauker, C. (2011). *Words and Images*. Oxford: Oxford University Press.

Gauthier, I., Tarr, M.J., Anderson A.W., Skudlarski, P. & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise recognizing novel objects. *Nature Neuroscience*, 2, 568-573.

Goldstone, R. L., Gerganov, A., Landy, D., & Roberts, M. E. (2008). Learning to see and conceive. In L. Tommasi, M. Peterson, & L. Nadel (Eds.) *The New Cognitive Sciences* (pp. 163-188). Cambridge, MA.: MIT Press.

Henderson, J. M., & Hollingworth, A. (1999). High level scene perception. *Annual Review of Psychology*, 50, 243-271.

Humphreys, G. W., & Riddoch, M. J. (1987). *To see but not to see: A case study of visual agnosia*. Hillsdale, NJ: Erlbaum.

Humphreys, G. W., & Riddoch, M. J. (2006). Features, objects, action: The cognitive neuropsychology of visual object processing, 1984-2004. *Cognitive Neuropsychology*, 23, 156-183.

Humphreys, G. W., & Riddoch, M. J. (2007). How to define an object. *Mind and Language*, 22, 534-547.

Hurley, S. (2008). The Shared Circuits Model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behavioral and Brain Sciences*, 31, 1-22.

Jacob, P., & Jeannerod, M. (2003). *Ways of Seeing: The Scope and Limits of Visual Cognition*. Oxford: Oxford University Press.

Kamiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge: MIT Press.

Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10, 494–501.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neuroscience*, 23, 571-579.

Landy, D., Goldstone, R. L.(2005). How we learn about things we don't already understand. *Journal of Experimental and Theoretical Artificial Intelligence*, 17, 343-369.

Lyons, W. (2008). *Perception and Basic Beliefs*. Oxford: Oxford University Press.

Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition*, 104, 19-46.

Mahon, B., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology, Paris*, 102, 59-70.

Mandler, J. (2004). *The Foundations of Mind*. Oxford: Oxford University Press.

Mandler, J. M., & McDonough, L. (1993). Concept formation in infancy. *Cognitive Development*, 8, 291-318.

Mandler, J. M., & McDonough, L. (1998). On developing a knowledge base in infancy. *Developmental Psychology*, 34, 1274-1288.

Matthen, M. (2005). *Seeing, Doing, and Knowing*. Oxford: Oxford University Press.

Milner, A. D., & Goodale, M. A. (2006). *The Visual Brain in Action* (2nd ed.). Oxford: Oxford University Press.

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

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

Pessig, J. J., & Tarr, M. J. (2007). Visual object recognition: Do we know more now than we did 20 years ago? *Annual Review of Psychology*, 58, 75-96.

Prinz, J. (2002). *Furnishing the Mind*. Cambridge: MIT Press.

Prinz, J. (2005) Beyond appearances: The content of sensation and perception. In T. Gendler & J. Hawthorne (Eds.), *Perceptual Experience* (pp. 434-460). Oxford: Oxford University Press.

Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, 9, 129-154.

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, 80, 127-158.

Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, 22, 463-475.

Rakison, D. H. (2005). The perceptual to conceptual shift in infancy and early childhood: A surface or deep distinction? In L. Gershkoff-Stowe & D. H. Rakison (Eds)., *Building Object Categories in Developmental Time* (pp. 131-158). Mahwah, NJ: Lawrence Erlbaum.

Rakison, D. H., & Butterworth, G. E. (1998a). Infants' use of object parts in early categorization. *Developmental Psychology*, 34, 49-62.

Rakison, D. H., & Butterworth, G. E. (1998b). Infants' attention to object structure in early categorization. *Developmental Psychology*, 34, 1310-1325.

Recanati, F. (2002). The Fodorian fallacy. *Analysis*, 62, 285-89.

Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience Supplement*, 3, 1199-1204.

Rives, B. (2010). Concepts and perceptual belief: How (not) to defend recognitional concepts. *Acta Analytica*, 25, 369-391.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and Categorization* (pp. 27-48). Hillsdale, N.J.: Erlbaum Associates.

Rosch, E., Mervis, C. B., Gery, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology,* 8, 382-439.

Roskies, A. (2008). A new argument for nonconceptual content. *Philosophy and Phenomenological Research*, 76, 633-659.

Scholl, B. (2001). Objects and attention: The state of the art. *Cognition*, 80, 1-46.

Schütz-Bosbach, S., & Prinz, W. (2007). Perceptual resonance: Action-induced modulation of perception. *Trends in Cognitive Sciences*, 11, 349-355.

Schyns, P. (1998). Diagnostic recognition: Task constraints, object information, and their interactions. *Cognition*, 67, 147-179.

Schyns, P., Goldstone, R., & Thibault, J. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, 21, 1-54.

Segar, C. A., & Miller, E. K. (2010). Category learning in the brain. *Annual Review of Neuroscience*, 33, 203-219.

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

Tanaka, J. W. (2001). The entry point of face recognition: Evidence for face expertise. *Journal of Experimental Psychology: General*, 130, 534-543.

Tanaka, J. W., & Gauthier, I., (1997). Expertise in object and face recognition. *The Psychology of Learning and Motivation*, 36, 83-125.

Tarr, M. J. (2003). Visual object recognition: Can a single mechanism suffice? In M. A. Peterson & G. Rhodes (Eds.), *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes* (pp. 177-211). Oxford: Oxford University Press.

Tye, M. (2005). Nonconceptual content, richness, and fineness of grain. In T. Gendler & J. Hawthorne (Eds)., *Perceptual Experience* (pp. 504-530). Oxford: Oxford University Press.

Weiskopf, D. A. (2007). Concept empiricism and the vehicles of thought. *Journal of Consciousness Studies*, 14, 156–183.

Weiskopf, D. A. (2009a). The plurality of concepts. *Synthese*, 169, 145-173.

Weiskopf, D. A. (2009b). Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research*, 79, 130–162.

Weiskopf, D. A. (forthcoming). The architecture of higher thought. In M. Sprevak and J. Kallestrup (Eds.), *New Waves in Philosophy of Mind*. New York: Palgrave Macmillan.