Perception, cognition, and perceptual categorization: the three systems account

ABSTRACT

The relationship between perception and cognition is a matter of intense controversy, both in philosophy and cognitive science. A prominent claim in recent work by Block (2014) and Burge (2010) is that there is a 'joint in nature' between perception and cognition. In this paper, I present a challenge to that view in the form of perceptual categorization, a form of processing that is not readily understood in strictly perceptual or strictly cognitive terms. Instead, I propose that perceptual categorization be understood as a further psychological kind in its own right, a position I term the *three systems* view. I develop an exploratory account of this position, and attempt to motivate it on the basis of both theoretical and empirical considerations. I conclude by suggesting that the three systems view may offer a useful framework in debates about the architecture of the mind.

1. Introduction

The relationship between perception and cognition has long been a source of controversy and in philosophy and psychology, with debates dating back at least to Plato's *Republic* (602c-e). In more recent years, the idea that perception and cognition are distinct has been championed by philosophers such as Dretske (1981), Fodor (1983), Burge (2010), and Block (2014).

In this paper, I will argue that those theorists sympathetic to a perception-cognition distinction face a challenge from a kind of processing at the intersection of the two that is sometimes called *perceptual categorization*. By this term, I mean the fast, automatic categorization of perceived objects in respect of learned categories.[1] One response to this challenge may to be abandon the perception-cognition distinction entirely. However, I will go on to suggest an alternative possibility, which I call the *three systems* approach. In short, this suggests that, insofar as we are inclined to think there is a joint in nature between perception and cognition, we should also be open to the possibility that perceptual categorization constitutes a *third* distinctive process.

The paper proceeds as follows. I begin in Section 2 by giving a rough summary of how contemporary defenders draw the distinction between perception and cognition, a position I term the *two systems* view. Next, in Sections 3 and 4, I introduce the idea of perceptual categorization, and go on to offer some challenges to the idea that perceptual categorization can be understood just in terms of either perception or cognition alone. In Section 5, I lay out an alternative proposal in the form of the *three systems view*, which holds that perceptual categorization can fruitfully be considered a distinct stage of processing intermediate between perception and cognition. Finally, in Section 6, I evaluate the three systems view in more detail, considering some objections and offering responses.

2. The two systems view

There are at least superficially important differences in the roles that perception and cognitive processes like thought and judgment play in the mind. Perception operates rapidly and automatically, whereas many cognitive processes such as reasoning typically operate slowly,

[1] In referring to this process as *perceptual* categorization, I do not assume that it is a function of perception proper. However, some such qualifying term is required, given the ubiquity of processes involving categorization in thought, mental imagery, and other contexts that are more clearly removed from perceptual inputs.

and are at least potentially under voluntary control. The short-term memory mechanisms often thought to underpin perception have a large capacity (Sperling, 1960; Block, 2007), whereas our ability to perform many cognitive tasks is strictly limited by the capacity of working memory (Cowan, 2001). Likewise, perception and cognition seem to differ in their representational repertoire. Even if one claims, for example, that we can perceive natural kinds or causal interactions (Siegel, 2010; Bayne, 2009), we surely do not perceptually represent highly abstract theoretical kinds like Gini coefficients or haloclines, nor complex logical structures like conditionals. By contrast, we have no difficulties in thinking or forming desires about these things. Conversely, it has sometimes been suggested (by, e.g., Peacocke, 1992) that perception has a finer grain than cognition, allowing us to perceive far more shades of red, for example, than we can form thoughts about.

Such differences between perception and cognition are often connected with the claim that they possess distinct representational formats. Dretske (1981) and Tye (1995), for example, take it to be partly constitutive of perception that its content has an *analogue* format, unlike the 'digital' or discrete contents of thought. Similarly, Block (2014) and Burge (2010) claim perception involves nonconceptual content and a sensory or iconic format, whereas thought is conceptual and propositional.

In addition to these more theoretical considerations supporting some kind of distinction between perception and cognition, there is also considerable phenomenological appeal to the idea. Perceptual and other sensory forms of experience are famously characterized by their 'qualia', the phenomenal characters associated with experience of qualities such as redness, pain, and saltiness. By contrast, it remains controversial whether canonical forms of cognition such as understanding and reflective thought have any phenomenal character at all; even those sympathetic to the idea that there is a phenomenology of thought must allow that its subjective character is sharply different from that of seeing and hearing, for example.

A final important difference sometimes claimed to exist between perception and cognition is that the former is in some sense *encapsulated* from the latter (Fodor, 1983); that is to say, the contents of perception are not amenable to correction or direct modulation by cognitive processes. What we see, in other words, is largely independent of what we believe. If I am looking at an illusion like the famous Muller-Lyer illustration, I cannot bring myself to *see* the lines as equal in length, even though I know they are. This is in contrast to the case of most cognitive states, which routinely interact with one another in ways that are sensitive to their respective contents; my belief that my friend Petra is in her office, for example, may be readily corrected by my remembering that she had earlier told me she was ill at home in bed.[2]

In light of these numerous prima facie differences between perception and cognition, it is not surprising that many philosophers and psychologists take the distinction between perception and cognition to be central to our understanding of the mind. While there are a great different ways to spell out the exact difference between the two, it is nonetheless

[2] It should be stressed this way of distinguishing between perception and cognition remains highly controversial. A considerable empirical and philosophical literature argues for the existence of top-down effects on perception (see, e.g., Vetter & Newen, 2014), while recent work on belief fragmentation suggests that certain cognitive states may be encapsulated from one another.

possible to give a highly generalized and schematic illustration of what these accounts have in common (Fig.1), and I will give them the label *two systems* views.[3] What I take to be essential to and common among these approaches is the idea that perception and thought differ in psychological role, phenomenology, content, or some combination of the above, and should therefore be considered distinct *psychological kinds*.

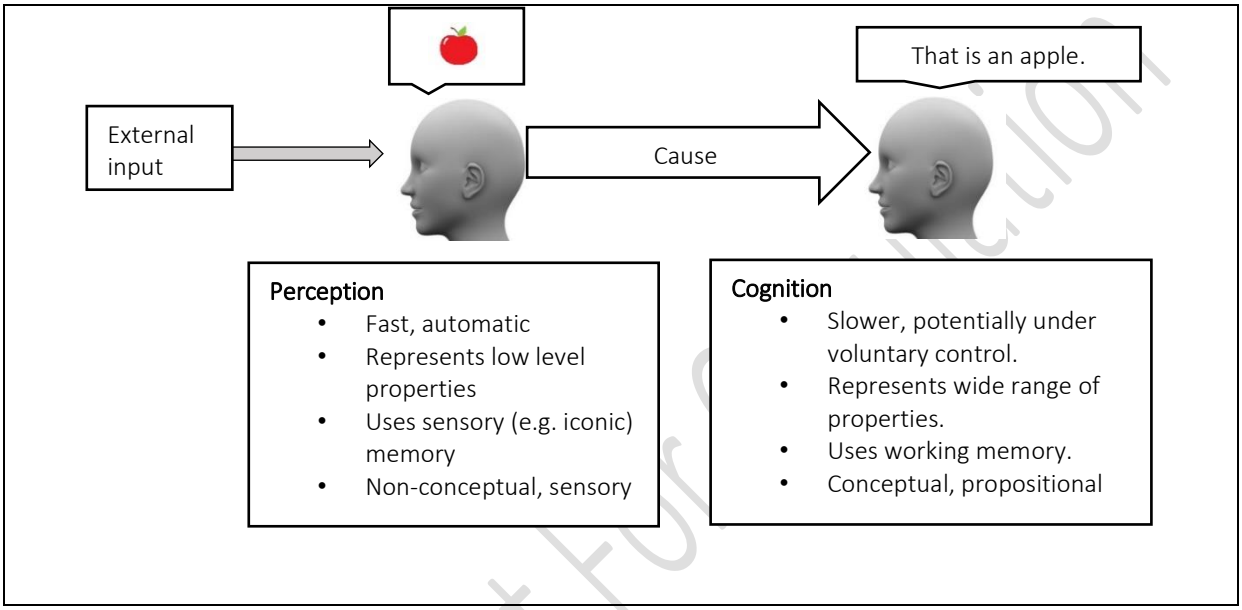

Fig. 1. A schematic illustration of the 'two systems' view

The two systems view in its various guises remains highly controversial, in spite of its theoretical and intuitive appeal. It has been challenged, for example, by predictive coding accounts of the mind championed by theorists such as Gary Lupyan (2015) and Andy Clark (2013). Additionally, the growing philosophical literature concerning alleged 'top-down' effects on perception from cognition has threaten to blur the boundary between the two processes.[4]

In this paper, I wish to address a different kind of challenge to two systems theorists. Instead of positing one too many systems, as predictive coding theorists might claim, I will explore the possibility that the position admits one too few. Specifically, I will claim that it should be expanded to include a third process in the form of *perceptual categorization*, as I will now describe.

3. Perceptual categorization

To sketch what I mean by perceptual categorization, we can begin by noting that perceptual experience makes us aware of different kinds of properties of objects and scenes. Some of

[3] The term 'two systems' view is also sometimes used to describe theories such as that developed by Kahneman (2014) that distinguishes between fast, automatic, non-conscious processes on the one hand and slower, reflective, conscious processes on the other. While there is no direct equivalence to be drawn between two systems views in the sense described here and that advocated by Kahneman, there are some important similarities. In particular, Kahneman regards perception as an example of a System 1 process, distinct from purely cognitive processes System 2 processes such as reflective reasoning.

[4] However, it should be noted that the mere existence of top-down effects would not count as decisive evidence against the two systems view (see Block, 2016).

these are the typical qualitative properties – color, size, shape, motion – and so on that are canonically associated with perception. Thus, as Prinz puts it, "[i]n vision, we experience stimuli as bounded wholes from a specific vantage point, occupying a specific size and position within the visual field" (Prinz, 2011). However, perceptual experience also arguably makes us aware of the *categorical* identity of objects. Thus, imagine you are walking down the street one day when you see your friend James. Your visual experience indeed makes you aware of James as a bounded whole of a certain size standing in a particular perspectival relationship to you. However, it arguably also serves makes you aware of him *as James* and *as a person*.

This latter form of awareness is what is commonly termed perceptual categorization or perceiving-as. It can be at least in principle distinguished from perception in the narrow sense of the term insofar as it involves the application of particular learned categorical identities (such as "person" or "James") to what is seen. Perception, by contrast, at least by the lights of many theorists (including, with qualifications, Burge and Block), encodes objects just in respect of more basic properties like those mentioned earlier. Perceptual categorization can also again be at least in principle distinguished from simply forming thoughts about what we see. Thus, looking at my friend James I may judge that he has taken the day off work, and that he is probably walking to the supermarket. However, it does not seem appropriate to say that I *see* him as having these properties, at least in the same sense that I see him as *being James*.

That said, in individual cases, it may be unclear whether seeing an object as possessing some property is best understood as a case of perceptual categorization or perception proper. This is especially true given that theorists such as Burge and Fodor (2015) take all perceiving to involve attributing properties to objects (as Fodor puts it, all perceiving is perceiving-as). I wish to leave it open, then, whether events like perceiving something as a discrete object, as solid, or as heavy are best understood as acts of perceptual categorization or as a more basic function of perception proper. Even granting that there are some obscure cases like this, however, we can nonetheless – at least for present purposes – draw a fairly robust distinction between our limited and largely innate capacity to represent sensory properties (hue, motion, pitch, and so on) and our capacity to represent objects in respect of learned high-level categories, including social and natural kinds like "picnic" or "pine tree".

4. Perceptual categorization: perception, cognition, or something else?

With this characterization in hand, we can now note that it is not immediately clear where perceptual categorization falls on the two systems framework illustrated above. While it shares some features with perception (in being fast and automatic, for example), it also seemingly involves high-level contents (like 'Barack Obama' or 'owl') of the form normally associated with cognition. We may thus wonder, with Bayne & Montague (2011), whether "some instances of 'perceiving-as' ought not be regarded as purely perceptual but are intermediate (or perhaps 'hybrid') states that straddle the divide between thought and perception."

I will shortly pursue this suggestion, arguing that perceptual categorization is different enough from perception and cognition that it might be fruitfully considered part of a third system. Before giving this account, however, it is worth exploring two alternative

possibilities: first, that perceptual categorization is a form of rapid automatic judgment directed at perceptual contents, and thus a form of cognition proper; and second, that is simply a form of high-level perception involving the same core mechanisms as those involved in perceiving color, shape, and so on.

Consider first the hypothesis that perceptual categorization is a form of cognition. A natural suggestion here would be that, when I see James, I form a rapid propositional *judgment* or *belief* (Byrne, 2009; see also Pitcher, 1971) to the effect that the person in front of me is James, with this judgment in turn underwriting my ability to recognize James. However, note that if such judgments are responsible for perceptual categorization, they must differ markedly from canonical cases of judgment and belief, such as judging that it will rain tomorrow. For one, perceptual categorization does not seem amenable to correction in light of my background beliefs. This as illustrated by pareidolic illusions, like the images shown below. Even if we are informed and confidently believe that what we are seeing, for example, a pool of soap bubbles, we cannot help but perceptually categorize it as an eye.[5]

Fig. 2. Examples of Pareidolia: an eye in soap bubbles and an owl in a coffee cup.

This provides some initial basis for doubting that perceptual categorization is a form of judgment: whereas our judgments are typically sensitive (to at least some degree) to other information at our disposal, no matter how firmly we believe that what we see is simply bubbles in a sink, we are seemingly irresistibly led to see the image as depicting an eye.

This is not intended as a decisive objection; there are, after all, some cases where cognitive states are seemingly insensitive to correction in light of new information, thus leaving open the possibility that perceptual categorization is a form of automatic judgment or belief. However, it should be noted that it would be quite unlike the familiar roster of cases where cognitive states are not amenable to fast correction. Implicit biases, motivated reasoning, and simple forgetfulness, for example, might lead someone to hold contradictory beliefs at different times or fail to update their judgments in light of new information. Pareidolia, however, presents a more difficult case, however, insofar as I may *simultaneously* have two occurrent representations of the image above with contradictory contents, judging "that is not an eye" while continuing to see it as an eye. Similarly, whereas most of our beliefs are at least in principle open to correction, it is hard to imagine what strictly *cognitive* shift could ever lead someone to stop seeing the image on the left as an eye.

[5] Note that unlike simpler illusory cases like the Muller-Lyer illusion, pareidolia need not involve any misrepresentation of low-level properties like color or size, but only the categorical identity of what is seen. Pareidolia, then, constitutes a distinctive kind of *categorical* known illusion case.

A further and perhaps more fundamental challenge for strictly cognitive accounts of perceptual categorization comes from the fact that many non-human animals, including invertebrates, are seemingly capable of it. For example, note that pigeons can be trained to discriminate novel pictures based on whether they contain images of people, and also seemingly learn to discriminate novel beach scenes from scenes of mountains or streets (Kirkpatrick et al., 2014). Rats, too, can learn to make same-different discriminations on novel stimuli, and can learn to distinguish novel pictures of chairs, flowers, cars, and humans on seemingly categorical grounds (Brooks et al. 2013). Perhaps most strikingly of all, bees can be trained to make same/different and above/below distinctions for wholly novel stimuli (see Chittka & Jensen, 2011 for a review).

Again, this objection is not decisive: it is possible that all these creatures possess propositional attitudes (Carruthers, 2009). However, this is far from the dominant view concerning the capacities of animals, especially simpler creatures such as bees, and is specifically rejected by many leading theorists (notably Burge, 2010). Another response on behalf of a cognitive account of perceptual categorization might claim that the kinds of apparent categorization performed by many or all the animals mentioned above differ from the human faculty in relying on simpler non-cognitive mechanisms. While this is a live possibility, there is evidence from evolutionary biology that the basic capacity for learning, memory, and recognition is shared across all vertebrates at least via the hippocampus and homologous structures, which might be taken as evidence of a shared basis for such discriminatory behaviour (Feinberg & Mallatt, 2016: Ch.6). More to the point, however, there seems to be a tension in claiming that, on the one hand, the sophisticated perceptual categorization behaviours of non-human animals can be explained in terms of simple non-cognitive mechanisms, while also insisting that at least superficially similar human capacities must involve cognition proper. More specifically, one might well think that to the extent that we are able to explain perceptual categorization in non-human animals without appeal to propositional attitudes, we might be able to – and should – do the same for human beings.

Setting aside for the moment, then, the idea that perceptual categorization is a form of cognitive judgment, let us examine the possibility that it is instead a form of perception proper. One such view (developed by Burge) might claim that when a creature becomes perceptually aware of a given animal as, say, a predator, it does so using a kind of non-conceptual attributive. A similar account for cases of perceptual categorization in humans is offered by Block (2014) who argues on the basis of experimental data that there is genuinely perceptual (that is, sensory and non-conceptual) representation of some fairly high-level properties such as emotion and facial expression.

I will not attempt to answer Burge and Block's specific arguments here. However, I would suggest that there are difficult questions that face any account that attempts to explain perceptual categorization in terms of processes with strictly sensory formats. Note in particular that Block (2015) suggests that the sensory representations of perception possess a *pictorial* mode of representation. That is not to suggest that mental icons for Block are like internal photographs. A better comparison, he suggests, are the representational forms open to painters: "an impressionist painter might represent a hand in broad brush strokes that do not explicitly represent the number of fingers or whether

one of them has a ring".

Even granting to Block, however, that sensory representations can possess at a high degree of *indeterminacy* of content, it seems hard to see how their representational range can capture all of the properties seemingly involved in perceptual categorization. That is to say, the pictorial features of representations involved in perceptual categorization in many cases arguably underdetermine their total representational content. Such representations may instead rely on previous acquaintance with appropriate stimuli as well as corresponding contexts and conventions in order to convey their content. Consider, for example, the cartoon below.

Fig. 3. A cartoon of a drunk.

I take it that this constitutes at least a prima facie instance of perceiving someone *as* drunk. However, the distinctively *pictorial* features of the image– the man's red nose, his bulging eye, the glass in his hand – are also compatible with many other interpretations. Someone with a different history of perceptual learning might, for example, see the man as *ill*, associating his red nose with his having a bad cold and his glass as containing medicine. Additionally, note that our interpretation of the man in this picture as drunk is almost certainly aided by the intrinsically meaningless symbols or 'emanata' above his head that conventionally depict inebriation (see below).

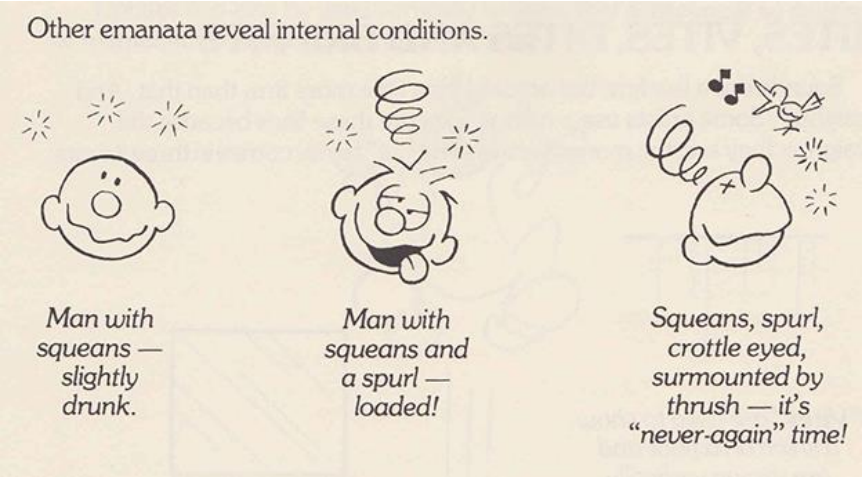

Fig. 4. The role of 'emanata' in depicting drunkenness. From the Lexicon of Comicana (Walker, 1980).

I do not wish over-intellectualize the process of understanding cartoons like the ones above: we are certainly not aware of making any conscious inferences about the content of Figure 3, and we do not consciously *use* symbols like the emanata in Figure 4 in a reflective manner (indeed, it is easy to fail to consciously notice them entirely). Nonetheless, thanks to most readers' familiarity with the relevant representational forms and conventions, these cues enable us to effortlessly represent the artists' desired content, in spite of the fact that this content is underdetermined by the purely *pictorial* aspects of the image in question.[6] This suggests that, to the extent that properties like 'looking drunk' feature among the contents of perceptual categorization, they rely at least in part on elements whose semantic contents go beyond merely pictorial forms of representation.

Of course, Block may simply deny that there is perceptual-level representation of properties like 'being drunk'. I am certainly open to the possibility (as mentioned earlier) that certain simple cases of apparent perceptual categorization like seeing something as an object or perhaps as predator or prey may ultimately be understood in terms of strictly perceptual. However, this leaves us wanting for an explanation as to how other forms of perceptual categorization, like seeing a person as your friend James or an object as a car, can be accomplished.

A quite separate consideration that might tell against a strictly perceptual view of perceptual categorization comes from the apparent finding that our capacity for categorical representation is dissociable from our capacity for representation for canonically perceptual properties such as color, shape, and size. This is illustrated by patients with associative agnosia, who are able to perceive (and accurately copy) images in respect of color, size, and shape, yet are unable to recognize what they are seeing. This suggests that there is more to seeing something as a given object than simply perceiving its shape, size,

[6] Such contextual and conventional cues may also help explain our ability to effortlessly recognize very abstract cartoons and caricatures of famous people.

color, and so on.[7]

Perhaps even more tellingly, many associative agnosics have category-specific deficits, whereby they are unable to recognize objects of a specific kind, such as animate objects in particular, suggesting that the impairment specifically affects areas of the brain involved in categorization, as opposed to, for example, areas involved in representations of complex object structure. Of course, this neurological data does not directly address the more philosophical question of whether perceptual categorization should be considered a species of perception proper, but does suggest that the underlying mechanisms for category recognition are somewhat distinct from those involved in awareness of color, shape, and size.

5. Perceptual categorization as a third system

These considerations are not intended as conclusive, and there are many ways to respond to them. The uneasy mix of perceptual and cognitive features involved in perceptual categorization might, for example, lead one to doubt whether there is any a clear joint between perception and cognition.

However, an alternative option for theorists inclined to defend a joint in nature between perception and cognition would be to consider whether perceptual categorization might be a further psychological kind in its own right. With this in mind, I suggest that we might best accommodate perceptual categorization within our account of the architecture of the mind by thinking of it as a process that serves as an *intermediary* between perception and cognition, whose function is, roughly speaking, to classify the outputs of perception in terms of both innately specified and learned categories relevant to an organism's goals.

This is the basic idea underpinning what I will call the *three systems* view. Note that this third system of perceptual categorization need not involve propositional attitudes, and perhaps not even concepts, but could perhaps function to create and assign simple semantic labels to different kinds of perceptual stimuli.[8] Thus one case might involve a non-human animal encoding a given sensory input in respect of ecologically relevant categories, such as "food", "predator", or "conspecific". More complex cases might involve seeing a given fruit as a banana, hearing a voice as sad, or recognizing a given person as James.

Before proceeding, I should address an immediate question that may arise, namely what exactly I am proposing in claiming that we should see the processes of perception, perceptual categorization, and cognition as three systems rather than two. In particular, one might worry that I am advocating for a mere notational variant of the two systems view; that, in proposing that perceptual categorization be treated as a distinct system, I am simply suggesting a kind of fairly unimportant taxonomic revision rather than a substantive theoretical change.

I would note that, even if one reads the view merely as proposing a shift in taxonomy,

[7] Note that the cases of pareidolia described earlier pose a similar problem for the sensory theorist as that presented by associative agnosia. Specifically, in both cases, it seems as though awareness of the sensory features of the object or image in question is accurate and veridical. What is illusory (in the case of pareidolia) or absent entirely (in the case of associative agnosia) is the categorical aspect of the visual experience. This creates a problem for any views that seeks to explain perceptual categorization just in terms of awareness of sensory features.

[8] See section 6, part (iii), below for a more extended discussion of representational format.

that does not mean that it is devoid of interesting content or potential utility. The way we carve up a territory on a map can make a great deal of difference to how we navigate the underlying terrain. A map that prints contour lines only every 20 metres change in height elides the different between bumpy undulating grasslands and a flat plain, whereas one that shows them every 10 metres may capture the differences. Similarly, one of the most the notorious taxonomic revisions in recent years, namely the reclassification of Pluto to a dwarf planet, was prompted by and continued to inform our revised view of the solar system as consisting not merely of large planets and small asteroids, but a range of bodies in between.

Nonetheless, I take the three systems view to be offering something more than a mere taxonomic division. Instead, I take it as proposing – tentatively – that perceptual categorization be considered a distinct *psychological kind*, distinct from both perception and cognition (see Fig. 5 below). By this, I mean that perception, perceptual categorization, and cognition form natural divisions in our architecture of the mind, underpinned by clusters of differences in respect of function, phenomenology, format, and content.[9] In other words, the distinction between perceptual categorization and perception is not arbitrary; it is not a matter of cutting a cake in various ways. Instead, there is, in the terms popularized by Ned Block (2014), a *joint in nature* between perceptual categorization and the other systems.

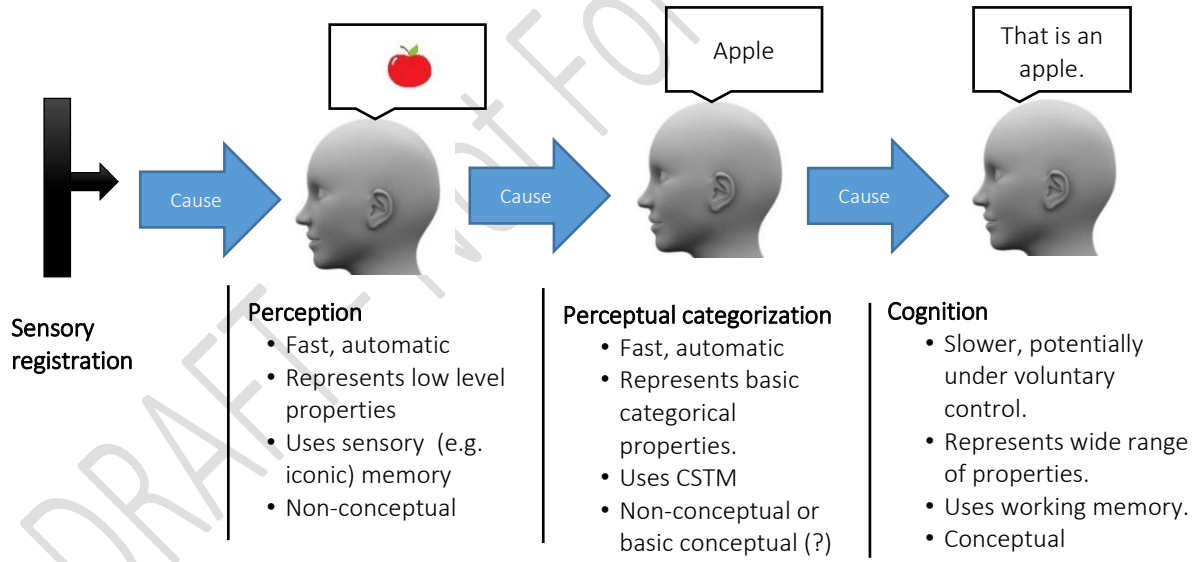

Fig. 5 – schematic illustration of the three systems view

Of course, one can allow that a given theory 'carves nature at the joints' without excluding the possible utility of other schema, and the specific frameworks that are relevant to a given enterprise will depend in part on one's interests: while the distinction between living and non-living systems may be crucial for a biologist, it may not be a distinction that has much relevance for astrophysics. However, within the relevant domain, we do well to carve nature at the joints insofar it allows for us to formulate generalizations with greater predictive and

[9] Note that, while I broadly agree with Block and Burge in taking perception to be a distinct psychological kind, I am more equivocal as to whether cognition is best considered a single kind or a set of fundamentally distinct processes (such as judgment and desire, for example) that interact in various ways. This question remains a difficult and outstanding problem for philosophy and cognitive science.

explanatory power.

Note that in arguing for the claim that perceptual categorization is a distinct psychological kind, this formulation of the three systems view voluntarily opens itself up to challenges on several fronts. For one, if it is found that distinguishing between perceptual categorization and other systems glosses over philosophically or scientifically important differences, or bogs us down with hair-splitting minutiae of classification, it may show itself to be failing to carve nature at the joints. Similarly, to the extent that we can get by with a simpler account, appealing simply to perception and cognition, it may thereby show itself to be theoretically profligate, failing to 'earn its explanatory keep'. With this in mind, my purpose in much of this paper is to press the claim that the three systems account can indeed earn its keep, insofar as it allows us to make sense of phenomena that sit uncomfortably within the traditional categories of perception and cognition.[10]

The above discussion should, I hope, provide a rough sketch of the proposal, and in a moment, I will spell out some arguments that I take to provide some tentative support for the model as well as clarify some of its features. Before doing so, however, I wish to draw attention to the fact that the three systems view in the bare form given thus far can remain neutral on a number of other questions. For example, I wish to take no firm stand on the issue of whether perceptual categorization involves *concepts*, or some more rudimentary form of categorical representation (however, see 6.1 below). Additionally, the three systems approach is compatible with a range of views concerning whether perceptual categorization is modular in the Fodorian sense (Fodor, 1983). While one might hold that the mechanisms of perceptual categorization are to a high degree encapsulated from central cognition, an alternate possibility would be that they are subject to modulation by occurrent cognitive states such as beliefs and desires. Finally, the view can remain fairly neutral on phenomenological issues, such as the rich content view (Siegel, 2010), and does not commit us to a position about what phenomenal character if any is evinced by the semantic representations processed in perceptual categorization.

5.1 – Broad theoretical considerations

I turn now, then, to some positive arguments for the three systems view, beginning with some broad theoretical considerations. I hope that in light of the arguments in the previous section, at least some of the initial appeal of this proposal should be fairly obvious. Specifically, it would allow us to accommodate the apparent ability of animals that may lack full-blown propositional thought to nonetheless use categories to structure their perceptual inputs. Indeed, it is easy to imagine how the ability to acquire novel categories and to deploy them to perceptual inputs might be of evolutionary value, even for creatures that lack propositional attitudes and possess relatively inflexible repertoires of behavior. Such an ability might allow organisms to rapidly identify novel signatures of predators, for example, without relying on

[10] There is a further, still stronger, reading of the three systems view, however, that I think is best avoided, namely that it singles out three *metaphysically* distinct kinds. I take it that the considerations that motivate a distinction between perception, cognition, and perceptual categorization follow from contingent features of human psychology rather than more fundamental reasons. I leave open the possibility, for example, that there might be creatures with mental architectures quite distinct from our own among whom we might find different clusters of properties at different levels of processing.

slower and riskier processes of associative learning.[11]

Additionally, the hypothesis that perceptual categorization is underpinned by a distinctive third system would explain some of the broad functional differences that distinguish it from canonical cases of perception and cognition. For one, it would provide a fairly straightforward account for associative agnosia, namely that the condition involved impairment of the third system, while leaving perception and cognition proper intact.

We might also suppose that the system underpinning perceptual categorization is fast and automatic, and that its outputs are at least somewhat resistant to correction by background beliefs. This would allow us to account for the existence of categorical forms of pareidolia problems, for example. Relatedly, but more speculatively, it might help us to account for the relative slowness of categorical forms of perceptual learning as compared to many forms of strictly cognitive learning. For example, I can learn via simply being told that a kingfisher is distinguished by its brilliant blue plumage and small beak, and thereby acquire corresponding beliefs about the appearance of kingfishers. However, actually acquiring the ability to reliably recognize kingfishers will typically require a longer period of learning via experience, specifically via seeing multiple instances of kingfishers.

5.2 – Evidence from short-term memory models

The above considerations hopefully provide some initial theoretical motivation for the three systems view. However, there is also some empirical reason to think that the three systems view may capture something important about the organization of the mind.

One promising line of research in this regard comes from emerging data in cognitive psychology concerning short-term memory mechanisms. Traditionally, debates about short-term memory have focused on nonconceptual sensory memory such as iconic memory and properly cognitive mechanisms such as working memory (Baddeley, 2003) and focal attention (Cowan, 2001). This has been reflected in philosophical treatments of perception and cognition, with a number of theorists (including Block, 2007, and Dretske, 1983) broadly identifying perception with sensory forms of memory and cognition with central cognitive processes such as working memory.

However, recent evidence suggests the existence of a third kind of short-term memory known as conceptual short-term memory or CSTM (Potter, 2012; Shevlin, 2017).[12] CSTM was initially posited by Mary Potter, and she characterizes it as "a mental buffer in which current stimuli and their associated concepts from long term memory… are represented briefly, allowing meaningful patterns or structures to be identified."

Though CSTM remains somewhat controversial, recent experiments suggest that it cannot be readily identified with working memory *or* sensory forms of memory such as iconic memory. In one crucial experiment, subjects saw 6-12 images presented in sequence

[11] Interestingly, similar evolutionary considerations for the structure of perceptual categorization in human beings have been advanced by psychologists in explaining the category-specific nature of associative agnosia. Thus Caramazza and Shelton (1998) suggest that "evolutionary pressures have resulted in specialized mechanisms for perceptually and conceptually distinguishing animate and inanimate kinds… leading to a categorical organization of this knowledge in the brain."

[12] Note that while Potter understands CSTM to be a properly *conceptual* mechanism, the results discussed below do not strictly require this. An alternate possibility to which I remain open is that information in CSTM in fact uses a nonconceptual but nonetheless categorical representational format.

for durations of 13, 27, 53, or 80ms (see Fig. 2). They were given a description of a target image (for example, 'wedding' or 'flowers') either 900ms or 200ms after seeing the images, and were asked to identify whether any images presented matched the description. On trials where the target stimulus was indeed present, subjects were given an immediate follow up recognition task to assess whether they could recognize the image when presented.

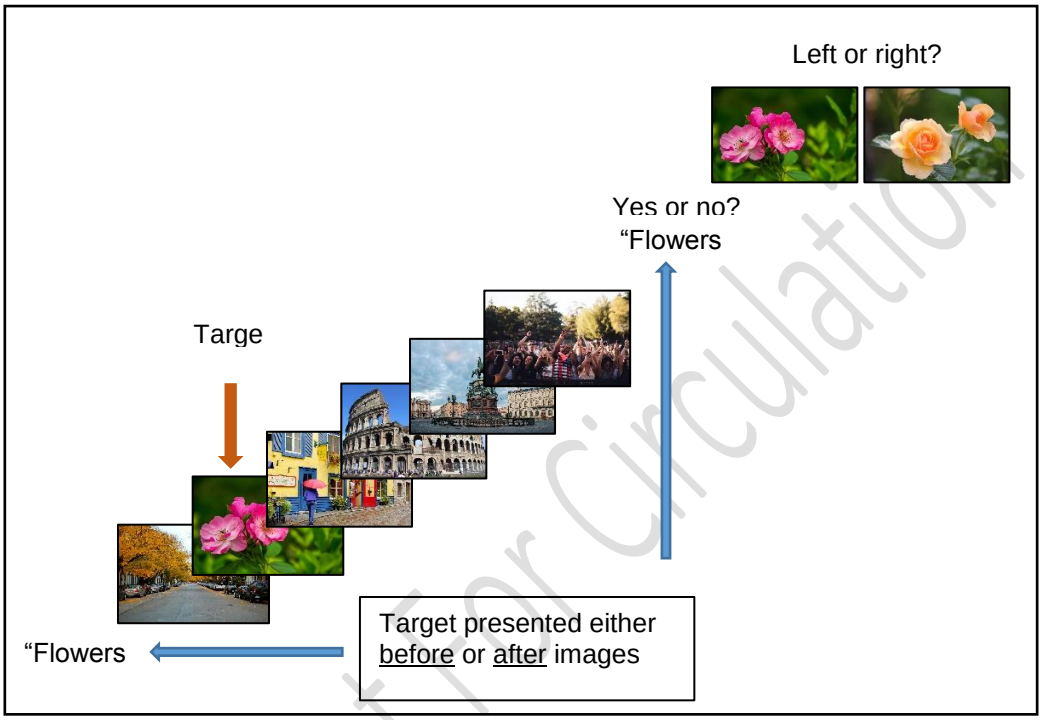

Fig. 6 (adapted from Potter et al. 2014). Sequences of 6-12 images were rapidly presented. Subjects were given a target description 900ms before or 200ms after presentation and instructed to say whether they saw a picture matching that description. In trials where a picture was presented, they performed a subsequent recognition task.

Subjects displayed above-chance performance on all measures, and in the follow up recognition task they were consistently able to recognize the target only if they had already detected its presence under the cued description. Crucially for present purposes, the results suggest that subjects were able to briefly retain categorical information about all or most of the 6-12 presented images. Otherwise, it is hard to see how they would be able to accurately report on the presence or absence of a given target when cued immediately *after* their presentation. This suggests that they were not merely storing the relevant information in working memory, given that this is generally taken to be a highly capacity limited mechanism (see, e.g., Cowan 2001). However, it is also unlikely that this information was stored in sensory memory: such buffers are normally taken to encode strictly sensory properties of stimuli such as color, size, and shape, and to be overwritten by sequential presentation of stimuli, whereas the tasks in this experiment seemingly required subjects to retain fairly high-level categorical information about images presented one after another.[13] These considerations lead Potter to claim that the results provide evidence for an intermediate high-

[13] There are of course a number of alternate debunking explanations for the results. For a more detailed discussion of these, see Shevlin (forthcoming).

capacity buffer in which the categorical identities of presented images are briefly stored, namely CSTM.

While CSTM remains somewhat controversial, I wish to emphasize that the putative features of CSTM match many of the features of perceptual categorization already described.[14] Specifically, it is fast, automatic, and has a brief duration, and serves to encode sensory information in respect of learned categories. This corresponds to the features of perceptual categorization already described. It may thus constitute a promising locus for perceptual categorization within the mind's different systems, and one that is plausibly distinct from the processes taken by many philosophers to underlie perception and cognition in the strict sense. This in turn suggests that the three systems model described earlier may potentially be able to be spelled out in terms of specific psychological mechanisms.

5.3 – Evidence from speed of category detection

A further strand of evidence for the three systems position comes from work that examines how rapidly subjects are able to access categorical information in presented stimuli. One sample experiment by Grill-Spector and Kanwisher (2005) examined subjects' ability to detect, categorize, and identify visual stimuli. All tasks involved presenting items for varied brief intervals (ranging from 17-200ms). The detection task required subjects to decide whether an object (as opposed to a texture) had been presented. The categorization task used a similar methodology but required subjects to press a button to indicate whether an item from a target category (e.g., car vs. not-car) had been displayed. Finally, the identification task required subjects to indicate whether the item belonged to a within-category class (e.g., German Shepherd vs. some other dog). Grill-Spector and Kanwisher discovered that while subjects were slower on the fine-grained identification task, their subjects' accuracy and speed was just as fast for the categorization task as it was for the detection task. Summarizing their results, they note that "detection and categorization performance require the same amount of information and processing time" and that "[b]y the time subjects knew an image contained an object at all, they already knew its category."

This result supports the intuitive idea that the basic categorical information is, at least in many cases, available for cognitive deployment from the first moment that a stimulus reaches awareness. In other words, it suggests that our awareness of categorical identity does not rely on making full-blown cognitive inferences that occur only after initial perceptual experience. This in turn might lend some tentative support to three systems model, insofar as it takes at least some forms of categorization of perceptual information to occur rapidly and automatically *prior* to the engagement of strictly cognitive mechanisms. Note also that subjects' relative slowness in fine-grained identification tasks as compared to basic categorization tasks also suggest that the relevant kind of categorization involves the application of relative generic semantic categories ("dog", "guitar", "car", and so on). This suggests that the intermediate process of rapid perceptual categorization is somewhat limited

[14] There is considerable further work in support of the CSTM hypothesis that I do not have space to discuss here. Two particularly interesting paradigms that explicitly invoke CSTM are (i) experiments that test subjects' ability to extract semantic information from complex arrays (Belke et al., 2008), and (ii) attentional blink paradigms that aim to detect subjects' representation of semantic mismatch in cases where they are not consciously aware of presented information (Luck, Vogel & Shapiro, 1996).

in its representational repertoire as compared to later cognitive processes, and marks a further functional distinction between perceptual categorization and cognition proper.

6. Evaluating the three systems view

The above arguments are not intended to show that the three systems approach is straightforwardly preferable to a two systems account that explains perceptual categorization just in terms of perception and/or cognition proper. However, I hope they give the view some initial plausibility. I now turn to consideration of some objections to the view, and in the process, hope to make clear both its core commitments and some of the different ways it can be cashed out.

6.1 – Representational format

An important initial worry for the three systems view concerns the specific format of representations involved in categorical perception. Thus far, I have merely suggested that they are not solely sensory, without committing to the claim that they are conceptual or propositional. Given that a number of theorists (including Dretske, Block, and Burge) explicitly appeal to representational format to distinguish perception and cognition, this may seem a worrying oversight, and one that threatens to undermine the view entirely. In particular, one might raise the concern that, if categorical perception involves representations with a propositional structure, it is *ipso facto* part of cognition; by contrast, if it involves non-conceptual representations derived from perceptual inputs, it is straightforwardly part of perception.

I would resist this move on two grounds. The first is that representational format is not the only basis on which to distinguish psychological kinds: considerations such as function, phenomenology, and psychological mechanism are also relevant. Thus I see no difficulty, for example, in someone's adopting the three systems view while simultaneously holding that the representations involved in categorical perception had propositional structure and content.

However, I think there may also be grounds for thinking that format itself can be used to further distinguish perceptual categorization from both cognition and perception proper. In particular, I am drawn to a view according to which the outputs of the mechanisms of perceptual categorization have a *hybrid* format, in which a semantic 'label' is indexed to an underlying sensory representation. I would suggest that we are all familiar with these kinds of hybrid representational formats in the form of labelled maps and diagrams (Camp, 2007). As a simple example, consider the seating chart shown below. Here, I can exploit the strictly symbolic elements of the representation to conclude, for example, that there are more women than men around the table. However, I can also use the analogue spatial structure of the chart to derive a very large number of more specific conclusions, for example that a man in the top-left hand corner is seated diagonally opposite another man, that the man to his right is sitting slightly further back, and so on.

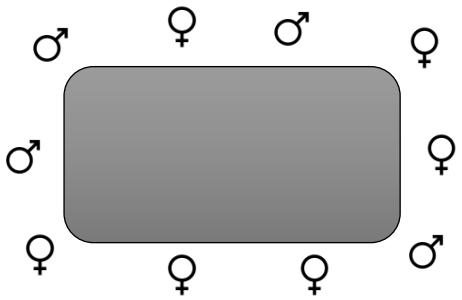

Fig. 7. A simple example of a hybrid representation showing men and women seated at various point around a table.

In asserting that the format of categorical perception involves a semantic element, I wish to leave open whether or not this element is strictly speaking *conceptual*. I am inclined to regard this as a matter that depends on one's broader theoretical commitments. At a minimum, I would suggest this semantic component is a repeatable, non-modality-specific attributive with a wide representational range encompassing both innate and learned categories. Additionally, unlike iconic representations, it does not need to rely whatsoever on structural isomorphisms with its representata in order to function representationally. However, this may not satisfy more demanding accounts of concepts. The semantic component in categorical perceptions may not, for example, be freely recombinable in a way that satisfies Evans' generality constraint (Evans, 1982). Additionally, I leave open the possibility that it may not be embedded within a full-blown propositional structure, but simply bound to an underlying sensory representation, with the result that it fail to count as a concept by the lights of Burge (2010). In any case, I am happy to leave this question open for present purposes.

The appeal to a hybrid format allows us to more clearly distinguish perceptual categorization from what I take to be the sensory format of perception and the full blown propositional structure common to thought. However, I would sugbgest that it is also somewhat independently motivated insofar as captures the phenomenological aspects of perceptual categorization, and the way in which our categorical awareness of objects is bound up with awareness of their lower level sensory properties. Thus, in viewing the pareidolic images in section (4) above, it is not the case that my categorical awareness of the eye or the owl is wholly removed from my sensory awareness of the image; while the eye and the owl are not simply determined by the sensory features of the image, the sensory features nonetheless contribute to them, affording them a definite shape, colour, and size.

Similarly, a hybrid account seems well placed to accommodate the experimental data of Potter et al. given earlier. Recall that subjects were only able to distinguish a presented image from another with the same rough semantic content if they had first detected it the relevant description. In other words, it seems that the process of categorizing the image appropriately facilitated a subsequent task (namely, recognition) in which the sensory features of the image were relevant. This is precisely what one might expect were categorical perception to involve a hybrid format, in which the representation of the semantic identity of a target also involved encoding at least some of its sensory features.

6.2 – Perceptual categorization and sensory biasing

A second objection to the account I have proposed thus far concerns the relation between perceptual categorization and a phenomenon (often referred to, somewhat unhelpfully, as *categorical perception*) whereby the perceived sensory qualities of a stimulus are influenced by the category according to which it is encoded. For example, in hearing phonemes with phonetic properties intermediate between /p/ and /b/, subjects typically report hearing strictly one or the other, rather than as intermediate sounds (Brandt & Rosen, 1980). Similar effects have been found in the case of colour, where subjects are faster as discriminating colours from different colour categories (like green and blue) than colours from a single category, even when the absolute degree of chromatic separation sizes are held constant (Gilbert et al., 2006).

One specific worry that comes from this research is that, in colour perception at least, such cross-category benefits are lateralized, occurring only on the right visual field; subjects display no similar comparative speed advantage when discriminating cross-category colours in their left visual field. The concern, then, is that this may suggest that perceptual categorization is properly *perceptual*, occurring early in experience prior to information from both visual fields being combined into a single image.

However, I think this objection mistakes the relation between perceptual categorization and biasing phenomenon such as these. While such biasing of sensory properties may be an *effect* of how one categorizes objects and stimuli, one that is lateralized to the right visual field (in adults, at least; see Franklin et al. 2008), I take it that perceptual categorization itself does not require it. Certainly, I do not believe categorization simply *consists* in such distortions at the sensory level, though it may promote such distortions in order to further aid category discrimination.

As an illustration of this kind of biasing at work, consider a recent experiment (Ester, Sprague, & Serences, 2017) in which subjects were taught during a one hour training session to assign fifteen different gratings to two unique categories depending on their orientation. Specifically, gratings that were oriented counterclockwise to a given boundary were assigned to the first category, and those oriented clockwise to that boundary to the second category. Following this initial training period, subjects were given a quite separate task involving letter identification during which a (task irrelevant) grating would appear in the background, and the responses of their visual areas to this stimulus were measured via fMRI and EEG.

The initial finding of the experiment was that subjects' early perceptual representations (in V1-V3) of the gratings were indeed biased by the categories they had just learned. Specifically, their representation of the gratings' orientation were biased *away* from their actual orientations and towards the centre point of the appropriate category.[15] Summarizing the result, Ester et al. note that "these results suggest that categorizing a stimulus alters how that stimulus is represented at the earliest stages of the visual processing hierarchy". A further finding of the experiment, particularly relevant for present purposes, is that the degree to which a stimulus was distorted away from its actual location depended on how close it was to a category boundary. Specifically, the closer a stimulus was to a boundary, the greater the degree of distortion, with entirely unambiguous stimuli being

[15] Note that, given the task irrelevant nature of the grating during the neuroimaging phase of the experiment, this biasing is hard to explain in terms of selective attention.

subject to no such distortion at all.

That our sensory systems can learn to 'round up' the sensory properties of ambiguous stimuli so as to be closer to the focal properties of their likely category is perhaps not surprising; a natural suggestion for why sensory biasing of this kind occurs at all is that is that it helps to reduce the risk of errors in categorical identification, and this risk is that much greater in the case of stimuli whose sensory properties would place them close to category boundaries. What is significant about this experiment for present purposes, however, is that it shows that, as suggested above, this biasing is a contingent effect of category learning, and is not essential to it: recall that categorically unambiguous stimuli were not subject to this biasing.

6.3 – Modality specificity

A final worry for the three systems view worth mentioning concerns how the proposed perceptual categorization system interfaces with the different sensory modalities. In particular, one might wonder whether there is strictly speaking *one* system of perceptual categorization, or many, corresponding to most if not all sensory modalities. If the latter, then one might ask whether it makes sense to speak of a *single* system of perceptual categorization at all.

Some initial grounds for this worry may come from the phenomenon of associative agnosia mentioned earlier. I suggested that the condition might be explained in terms of a deficit to the systems of perceptual categorization. However, associative agnosia is typically modality-specific, insofar as sufferers lose the ability to categorize objects visually or aurally rather than across different modalities. This might suggest that in fact, perceptual categorization is carried out separately across different sensory modalities.

The question of whether perceptual categorization relies on different cognitive mechanisms across different sensory modalities or instead involves a single multi-modal process is properly in the domain of cognitive neuroscience. Indeed, this is an area of active debate in the contemporary literature.[16] However, my hope is that I have framed the three systems view in such a way that it can remain neutral across these two alternatives. To begin with, it is hard to see how any account of perceptual categorization could deny that all relevant processing is amodal; categorizing something as a fish either by seeing a fish or hearing the word "fish" will involve matching quite different sensory inputs in the two cases, after all. Whether or not perceptual categorization as a whole is best viewed as a unified or disparate set of cognitive mechanisms will turn not on this question, but rather on, among other things, whether the various forms of perceptual categorization culminate in a distinct stage of amodal semantic processing.

Additionally, I would suggest that the three systems view is best understood as tracking relatively *abstract* psychological kinds, and is thus somewhat removed from more fine-grained questions about the precise cognitive and neural mechanisms involved. Just as it is no real objection to the view that perception itself is a single psychological kind that it is underpinned by different sensory mechanisms, it is compatible with the three systems view

[16] See, for example, Peelen & Caramazza (2012) for the claim that object recognition involves amodal processing the anterior temporal lobe. Bonner & Price (2013) by contrast argue that such processing involves some degree of modality-specificity.

that, at a more fine-grained level of explanation, we must appeal to an array of distinct cognitive mechanisms to explain the various forms of perceptual categorization.

(7) Conclusion

My primary goals in this paper have been twofold. First, I have sought to draw attention to the phenomenon of categorization and its uncomfortable position within the two systems framework widely defended in philosophy and cognitive science. Second, I have suggested that a compelling response to these difficulties is to expand our architecture of the mind to recognize a third distinct stage of processing intermediate between perception and cognition, and to regard this stage as a putative psychological kind in its own right. This proposal is tentative, of course, but I believe it may offer us a powerful and elegant framework for understanding the place of categorization within the mind.

An important alternative to this position that I have not considered is that both 'systems views' are flawed, and that the perception-cognition distinction itself should be reconsidered in favor of a much more fine-grained framework. This is the kind of position advanced by Clark (2013), for example, when he claims, for example, that "the lines between perception and cognition [are] fuzzy, perhaps even vanishing". While there is much about this position that appeals (for example, its ready ability to account for apparent cases of top-down effects on perception), whether it should be preferred to the account I have presented here will depend, in part, on its ability to account for the stark apparent differences between perception and cognition in psychological role, phenomenology, and format. As matters stand, I am of the belief that these differences do motivate a fundamental distinction between perception and cognition, and that the best way to accommodate many seemingly intermediate phenomena is instead to admit of a third distinct psychological kind intermediate between the two.

REFERENCES

Baddeley, A.D. (2003). Working memory: Looking back and looking forward .*Nature Reviews Neuroscience, 4*(10), 829-839.

Bayne, T., & Montague, M. (2011). Introduction to Cognitive Phenomenology. In Bayne, T., & Montague, M. (eds), *Cognitive Phenomenology*. Oxford University Press UK.

Belke, E., Humphreys, G., Watson, D., Meyer, A. and Telling, A., (2008). Top-down effects of semantic knowledge in visual search are modulated by cognitive but not perceptual load. *Perception & Psychophysics*, 70 8, 1444 – 1458.

Block, N. (2014). Seeing-As in the Light of Vision Science. *Philosophy and Phenomenological Research*, *89*(3), 560-572.

Block, N. (2015). Solely Generic Phenomenology. *Open MIND 2015*.

Block, N. (2007). Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience, *Behavioral and Brain Sciences* 30, pp. 481–499.

Block, N. (2016) Tweaking the Concepts of Perception and Cognition. *Behavioral and Brain Sciences*, 39: 21-2.

Bonner, M.F., & Price, A. R., (2013). Where Is the Anterior Temporal Lobe and What Does It Do? *Journal of Neuroscience* 6 March 2013, 33 (10) 4213-4215.

Brandt, J., & Rosen, J.J. (1980). Auditory phonemic perception in dyslexia: Categorical identification and discrimination of stop consonants." *Brain and Language*, vol. 9, no. 2: 324–337

Burge, T. (2010). *Origins of objectivity*. Oxford: Oxford University Press.

Byrne, A. (2009). Experience and content. *Philosophical Quarterly* 59 (236):429-451.

Camp, E. (2007). Thinking with maps. *Philosophical Perspectives* 21 (1):145–182.

Caramazza, A., & Shelton, J. R. (1998). Domain-Specific Knowledge Systems in the Brain: The Animate-Inanimate Distinction. *Journal of Cognitive Neuroscience,10*(1), 1-34.

Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R.W. Lurz (ed.), *The Philosophy of Animal Minds*. Cambridge University Press 89-107.

Chittka, L., & Jensen, K. (2011). Animal Cognition: Concepts from Apes to Bees. *Current Biology, 21*(3)

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36 (3):181-204.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences 24*1, 87-114.

Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press.

Ester, E.F., Sprague, T.C., & Serences, J.T. (2017). Category learning biases sensory representations in visual cortex. bioRxiv 170845; doi: https://doi.org/10.1101/170845

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

Feinberg, T.E., & Mallatt, J. (2017). *The ancient origins of consciousness: how the brain created experience*. The MIT Press.

Fodor, J. (1983). *Modularity of Mind*. MIT Press.

Franklin, A., Drivonikou, G.V., Bevis, L., Davies, I.R., Kay, P., & Regier, T. (2008). Categorical perception of color is lateralized to the right hemisphere in infants, but to the left hemisphere in adults. *Proceedings of the National Academy of Sciences* 105(9):3221-5.

Gilbert, A.L., Regier, T., Kay, P., & Ivry, R.B. (2006). Whorf hypothesis is supported in the right visual field but not the left. *Proceedings of the National Academy of Sciences* 103:489–494.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

Kirkpatrick, K., Bilton, T., Hansen, B. C., & Loschky, L. C. (2014). Scene gist categorization by pigeons. *Journal of Experimental Psychology: Animal Learning and Cognition, 40*(2), 162-177.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature, 383*(6601), 616-618.

Lupyan, G. (2015). Cognitive penetrability of perception in the age of prediction: Predictive systems are penetrable systems. *Review of Philosophy & Psychology* 6 (4):547-569.

Peelen MV, Caramazza A (2012) Conceptual object representations in human anterior temporal cortex. *Journal of Neuroscience* 32:15728–15736.

Pitcher, G. (1971). *Perception*, Princeton: Princeton University Press.

Potter, M.C., Wyble, B., Hagmann, C.E., & McCourt, E.S. (2014): Detecting meaning in RSVP at 13ms per picture. *Attention, Perception, & Psychophysics*, 762, 270-279.

Potter, M. C. (2012). Conceptual Short Term Memory in Perception and Thought. *Frontiers in Psychology*, 3:113.

Prinz, J. (2011). The Sensory Basis of Cognitive Phenomenology. In T. Bayne, & M. Montague, (eds.), *Cognitive Phenomenology*. OUP.

Shevlin, H. (2017). Conceptual Short-Term Memory: A Missing Part of the Mind? *Journal of Consciousness Studies*, Vol. 24, No. 7-8, July/August 2017.

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

Sperling, G. (1960). The Information Available in Brief Visual Presentations, *Psychological Monographs: General and Applied* 74, pp. 1–29.

Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.

Vetter, P., and Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness and Cognition.* 27, 62–75.

Walker, M. (1980). The *Lexicon of Comicana*. iUniverse.