

A role for the prefrontal cortex in supporting singular demonstrative reference

Felipe Nogueira de Carvalho (Universidade Federal de Minas Gerais)

Albert Newen (Ruhr-Universität, Bochum)

Author's pre-print. Full article now published in **Journal of Consciousness Studies**
26 (11-12): 133-156 (2019).

Abstract: One of the most pressing questions concerning singular demonstrative mental contents is what makes their content singular: that is to say, what makes it the case that individual objects are the representata of these mental states. Many philosophers have required sophisticated intellectual capacities for singular content to be possible, such as the possession of an elaborate scheme of space and time. A more recent reaction to this strategy proposes to account for singular content solely on the basis of empirical models of visual processing. We believe both sides make good points, and offer an intermediate way of looking into singular content. Our suggestion is that singular content may be traced to psychological capacities to form flexible, abstract representations in the prefrontal cortex. This allows them to be sustained for increasingly longer periods of time and extrapolated beyond the context of perception, thus going beyond low-level sensory representations while also falling short of more sophisticated intellectual abilities.

I – Introduction

When we perceive an object in front of us we are normally in a position to make a demonstrative reference to that object. Some philosophers have argued this is

the most fundamental form of reference there is, upon which all other reference rests (Strawson 1959). Others have assigned demonstrative reference a foundational role in language learning and the acquisition of empirical knowledge (Burge 1977, Brewer 1999).

But regardless of the undeniable importance of demonstrative reference, these are not the main questions that shall concern us in this paper. Rather, we wish to take the focus away from demonstrative reference *per se* and examine the underlying mental states that support particular acts of reference – states that we shall call ‘*singular demonstrative mental states*’, in order to capture two important elements of their contents: the first, singular element, concerns the question of what kind of entities are represented in the content of these mental states. To mark the content as ‘singular’ is to say that individual objects are the *representata* of these mental states, rather than properties, events, etc. These objects are the main subject matter of the content.

The second, demonstrative element, concerns the metasemantic question of how the content is determined. To mark the content as ‘demonstrative’ is to say that the intentional object of the content is determined through a contextual, perceptual relation to things in the subject’s external environment.

Classifying these mental states in this manner allows us to distinguish them, on the one hand, from singular contents that are not demonstrative (like the content of a thought about an individual that is not perceptually present), and on the other hand, from demonstrative contents that are not singular (like the content of a thought about a perceptually presented type of color). This allows us to demarcate very precisely our object of study, at the same time that it renders explicit two important questions about these mental states that we should be able to explain:

- **The demonstrative question:** what makes the contents of these mental states ‘demonstrative’, i.e., contextually connected to things in the external world?
- **The singular question:** what makes these contents singular, i.e., contents that represent individual objects as opposed to properties, events or some other suitable candidate for the *representata* of these mental states?

In this paper we will be mostly concerned with the second, singular question, although we will also touch on the demonstrative question. We will propose a way of answering the singular question that stands in a middle ground between two popular positions among theories on singular content: on the one hand, from philosophers who seek to account for this question on the basis of empirical models of visual processing borrowed from cognitive science, and on the other hand, from philosophers who postulate more intellectually sophisticated psychological abilities in order for singular content to be possible. Although the range of such abilities is undoubtedly wide, we will focus here on abilities that evolve around the possession of an elaborate conception of space and time, as required by Evans (1982), Quine (1995) and Hatfield (2009). For ease of exposition, we shall call them *perceptualist* and *intellectualist* theories respectively.

We believe both sides make good points that should be taken into account, and we offer a novel way of looking into singular mental content which attempts to do justice to both insights. We agree with perceptualist theories that we should be able to elucidate the notion of singular content at a more basic level of cognition, without requiring too many intellectual capacities on the part of the subject. Moreover, we

also agree that philosophical analysis should be informed by the best and most advanced cognitive science, and therefore that theories of singular content should go hand in hand with empirical theories of object representation in perception.

However, we also agree with intellectualism that when we come up with a theory of singular content, we want more than to know how visual systems are able to single out and keep track of discrete perceptual units in a visual array. We are also interested in the subject's psychological capacities to apprehend her external environment as structured into particular objects. And particular objects are more than discrete clusters of properties that our visual systems are sensitive to; they are also unique entities that persist in a time that transcends the subject's perceptual interactions with the object, and that move through a space that extrapolates the subject's visual field. And low-level visual processes posited by perceptualism tell us little about these psychological capacities. In this sense, to mark the content of a mental state as singular is to ascribe the subject a psychological capacity to apprehend her external environment as structured in this manner.

In this paper we propose a novel way of looking at the role of cognitive science in helping us account for singular demonstrative thought. In order to do so we will draw attention to a different body of empirical research, one that's been largely ignored by perceptualist theories of singular content. While perceptualist theories put an excessive focus on the perceptual mechanisms underlying visual representations of objects, we will suggest that a promising path to elucidate the singular question may lie in the role of the prefrontal cortex (PFC) in creating and sustaining more abstract object representations in the absence of perceptual input. Empirical research on object representations in the PFC can give us a clue about how our capacity to represent the continued existence of objects may be extrapolated from the context of perception. As

this does not require much else from the subject other than a capacity to interact with the object, represent it in a more abstract manner, and sustain that representation in mind across perceptual and non-perceptual contexts, it allows us to illuminate the singular question in a less intellectualist terms.

The structure of the paper is as follows. Section II will present two prominent contemporary perceptualist theory: John Campbell's attentional model (2002), and Joseph Levine's non-attentional model (2010). These theories are quite well known in the literature on singular content and provide excellent examples of what we call perceptualism, hence our focus on them. Section III articulates the intellectualist challenge to this type of theory, and sketches some proposals of how to elucidate the singular question within a more intellectualistic framework.

In section IV we examine two other intermediate level theories between perceptualism and intellectualism, namely, Imogen Dickie's and José Luiz Bermudez's, in comparison to our own account. Finally, in section V we will propose our own intermediate theory of singular demonstrative content. In our picture a subject may move beyond sensory representation of objects towards a more complex understanding of the spatiotemporal behavior of objects based on psychological capacities to sustain an object representation in mind while detached from the context of perception. These capacities are supported by abstract representations formed in the prefrontal cortex.

On this basis, we become able to represent the continued existence of the object beyond our perceptual interaction with it, thus giving substance to the idea that a subject in this situation is starting to grasp her external environment as structured into individual objects.

II – Perceptualist theories of singular content

The basic presupposition behind the perceptualist strategy is that there are natural perceptual processes and mechanisms that function to single out objects in the world, independently of concepts or descriptive material in the mind of the subject. The existence of these mechanisms is supported by empirical evidence from cognitive science. They are responsible for establishing a certain relation R between the subject and objects in the world, a relation that is supposed to fill two theoretical roles:

1. It provides the point of contact between the mind and the external world
2. It explains how individual objects are represented in the content of visual perceptual states

Both of these points seek to elucidate the demonstrative and the singular questions respectively. For once we have explained, on the basis of R , why individual objects are the *representata* of visual perceptual states caused by the workings of these mechanisms, we can then say that mental states based on the perception of objects – like the demonstrative mental states we’re analyzing – will ‘inherit’ the singular content of perception, and come to represent the very same objects represented in the content of the underlying perceptual states.¹

In this picture it is R what gives demonstrative mental states their singular content; singularity comes from below, so to speak, from the more primitive level of perception. The resulting picture would look something like this:

¹ The notion of ‘inheritance’ may be understood in Burge’s sense, when he talks of demonstrative applications in thought or language – like singling out an object with the demonstrative ‘that one’ – “inheriting”, or “taking over”, the referents of their counterpart perceptual applications. These perceptual applications are, precisely, the processes and mechanisms responsible for forming perceptual representations of objects in the visual system (see Burge 2010: 546).

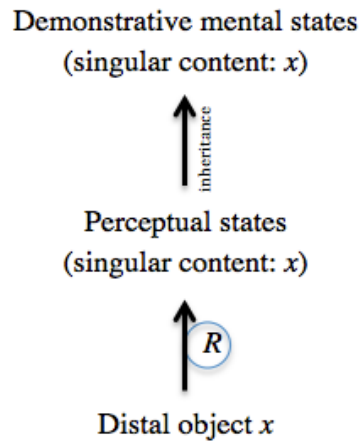


Fig.1

John Campbell was one of the first philosophers to systematically combine empirical models of perceptual processing with a philosophical account of perceptual demonstratives. According to Campbell, the perceptual relation that puts us in contact with external objects is an attentional relation, which fills two explanatory roles: it explains how objects are selected and processed by the visual system, and it enables us to perceive the object as ‘experientially highlighted’ (2002).

The first theoretical role is elucidated through Treisman & Gelade’s Feature Integration Theory (FIT), a model of perceptual processing where attention serves as the glue that binds various features together, enabling a perceiver to consciously experience these features as features of a single bounded object.² In the FIT model, simple visual features like color or orientation are first represented pre-attentively and in parallel in modular processing streams called feature maps. Conscious spatial attention comes in as the process responsible for binding together features detected in separate feature maps to one and the same object, which is none other than the object that occupies the attended location.

² See Treisman & Gelade 1980

As to the second explanatory role, Campbell argues that it is only in virtue of being in perceptual states with the phenomenal property of experiential highlighting that a subject is able to have singular demonstrative thoughts about external objects. This notion is first introduced in terms of the metaphor of “highlighting a portion of the visual field with a yellow marker” (Campbell 2002: 4), which seeks to capture the distinctive phenomenology of conscious visual attention. When a subject is presented with objects x and y but consciously attends to x , this act of conscious attention brings x into focus, at the same time that y is assigned to the background; x is then “highlighted” in her conscious experience of the scene, to the detriment of y . But if the subject then consciously shifts her attention to y , the reverse occurs, and y will be highlighted in her experience to the detriment of x .

Both of these elements are required for singular demonstrative thought to be possible, and both are explained through conscious attention. First of all, if your perceptual system has managed to process visual information about an object but you do not perceive this object as experientially highlighted, you will not be in a position to have singular demonstrative thoughts about it. To illustrate with an example, imagine you are on top of a mountain gazing distractedly at the forest landscape below. You are perceiving a lot of trees, but you do not zoom upon any one tree in particular. Your experience remains, as it were, as of ‘a sea of trees’. Although your visual system has nonetheless processed visual information about the trees, you are not in a position to think ‘that one’ of one of the trees in particular.

The situation changes when you consciously attend to a particular tree, and come to perceive it as experientially highlighted. When you think about ‘that tree’, this is the tree you are referring to: the one that is experientially highlighted in your experience, caused by your consciously attending to it. At the subpersonal level this

attentional act sets perceptual mechanisms in motion that will bind features found at the attended location to one and the same object, and at the personal level you will perceive the tree as experientially highlighted. When both of these conditions are in place, you are in a position to have a demonstrative thought about the tree. The tree you are thinking about is none other than the tree that is experientially highlighted in your experience, and whose features your perceptual system has bound together as features of one and the same object.

There are, however, some problems with Campbell's account. Regarding the second explanatory role of conscious attention, Wayne Wu (2011) has argued that the phenomenal property of experiential highlighting is more plausibly associated with some aspect of cognition, an element of the phenomenology of thinking demonstratively about an object. Hence, it cannot be what explains demonstrative thought in the first place. Regarding the first explanatory role, there is now evidence of pre-attentive perceptual mechanisms of object representation, which speak against *R* being an attentional relation. If this evidence is correct, it seems that once attention is brought into the picture objects have already been selected and represented by the visual system.

Joseph Levine explores some of this evidence in his account of singular demonstrative thought, based on Zenon Pylyshyn's (and collaborators') well-established work on visual indices, or FINSTS.³In Pylyshyn's model FINST's are mental particulars in the early vision system that are captured by external objects in a purely bottom-up manner, allowing the subject to keep track of up to four objects

³ See Pylyshyn 2001/2007. FINST's stand for Fingers of INSTantiation, a term designed to capture Pylyshyn's imaginative analogy with the superhero "plastic man", who is able to stick his plastic fingers to particular objects and keep them there as the objects move randomly through space (Pylyshyn 2007: 14)

during a period of observation without having to selectively attend to the location of every target to be tracked. The most important evidence comes from the Multiple Object Tracking (MOT) experimental paradigm, where subjects are able to keep track of up to four targets moving randomly across a screen among qualitatively identical non-targets. The high success rate of this task of about 85% (Pylyshyn 2007: 36), supports the existence of a pre-attentive mechanism of object selection and tracking, for subjects in the task cannot consciously attend to each target to be tracked.

Based on this evidence, Levine sketches a perceptual/cognitive architecture hierarchically structured in three levels. At the highest level we have demonstrative thoughts, whose constituents are mental representations of the form [that x]. These representations are structured as ‘mental pointers’, which point to perceptual representations of objects that have been previously selected by attentional mechanisms (in this case, a perceptual representation of the object x). But Levine takes these attentional mechanisms to operate not on external things directly but on pre-attentive perceptual representations of objects, which he identifies with Pylyshyn’s FINST’s. The function of these attentional mechanisms is to select a subset of these pre-attentive representations for further cognitive processing.

FINST’s are assigned to objects automatically and pre-attentively, independently of concepts or descriptive material in the mind of the thinker. It is thus a very good candidate for elucidating the demonstrative question. Moreover, Pylyshyn’s evolutionary story about how FINST’s get assigned to objects supports the claim that FINST’s capture ordinary material objects, since they are the entities with the right spatiotemporal properties and causal power to attract FINST’s in the first place. And if we stick to Levine’s three-level picture of perceptual and cognitive

processing, attentional processes operate directly on FINST representations, thus inheriting the singular content of pre-attentive perception.⁴

Finally, since mental demonstratives are nothing but internal pointers to these attentive perceptual representations, demonstrative mental states will also inherit the singular content of attentional perception. In this manner, both the demonstrative and the singular question are elucidated by the perceptual relation R , established when FINST's are automatically captured by objects in the world.

To conclude this section, we have examined two different options of what the fundamental perceptual relation R might be, namely, conscious attention and the assignment of visual indices (or FINST's). We have also seen that this relation R is taken to fill two theoretical roles, elucidating both the demonstrative and the singular questions. But how well do these theories manage to elucidate these two questions?

Regarding the demonstrative question, we believe both of these theories have the resources to answer it in a satisfactory manner. When it comes to this question we are more than happy to endorse an appeal to empirical models of visual processing in helping us adequately describe it. Which of these options will turn out to be the most adequate will depend on our best philosophical analysis coupled with the latest experimental evidence on how our visual system works, something we can leave open for the time being. As far as we are concerned, it is possible that one of these mechanisms, or perhaps a combination of these, may be used in an empirically responsible explanation of how our mental states make contact with external objects

⁴ There is a disagreement between Pylyshyn and Levine on whether FINST's qualify as perceptual representations or not; Levine thinks they do, while Pylyshyn argues they merely point to objects in the world without representing them. This debate, however, is largely terminological, and has no bearing on the present discussion.

via visual perception. We will, therefore, remain neutral on which of these perceptual mechanisms best characterizes the point of contact between mind and world.

We are, however, less optimistic about low-level perceptual processes when it comes to the singular question. As we will clarify in the next section, there are good reasons to suppose that something else, over and above low-level processes of perceptual tracking, is required for singular content to be possible. This will be the main subject matter of section III.

III – The intellectualist challenge to perceptualism

As we have seen in section II, there is empirical evidence that our visual system is able to single out discrete perceptual units in a complex visual array, which can then be visually tracked over a period of observation. Nevertheless, we may still legitimately ask: what makes it the case that the subject whose mental states we're analyzing represents particular objects in this case, rather than repeatable instances of a more general property? We are talking here of a well-known contrast between a creature who, when perceptually confronted with a particular apple, is only able to recognize that it is once again in a situation where a feature like 'appleness' is present in its external environment, *versus* a creature that is able to grasp that it is faced with a particular apple; an apple that might be numerically identical to an apple encountered in earlier occasions, and that might be encountered and re-identified again in the future.⁵ In the first case, the content of the creature's representational state would be captured in terms of a feature-placing structure like 'appleness', or

⁵ This contrast was originally drawn by Strawson (1959: 214ff). See also Bermudez (2003) and Cussins (1992) for informative discussions on the 'feature-placing' and the 'particular-involving' levels of experience.

‘appleness here again’, while in the second case we would attribute a singular content of the form ‘this particular apple’.

The problem in the context of the present discussion is that low-level processes of object perception and tracking fails to differentiate between these two scenarios, since in both cases the creature’s perceptual system would need to segment the visual scene into discrete perceptual units and track their movements in the visual field. What is needed, according to intellectualist accounts of singular demonstrative content, is a more sophisticated understanding of ‘objecthood’. For it is precisely this understanding what allows us to attribute singular content to the subject’s mental state, rather than a feature-placing content that merely captures her capacity to respond to recurrent instances of features. Otherwise we would be merely smuggling objects for free into the content of demonstrative mental states without warrant (Hatfield 2009: 215).

The point here is that there are important psychological dimensions to singular content that perceptualism leaves out. We ascribe singular content in order to mark (among other things) certain psychological capacities on the part of the subject. So the elements we assign to representational content, as well as the way in which they are combined, should reflect the way the subject grasps the structure of her external environment. In assigning singular demonstrative content to a subject’s mental state, for example, we are ascribing to this subject a capacity to apprehend her external environment as structured into individuals (object-situation representation), and not just repeatable instances of bounded properties (feature-situation representation): that the thing the subject is cognitively engaged with is a unique, persisting individual that will continue to exist after the encounter has ceased. But if the subject shows no signs of such psychological capacities, the ascription of singular content is explanatorily

idle. The capacity to perceive spatially differentiated spatial structures and perceptually track them through attention, for example, may be easily captured by a content structured in terms of repeatable instances of a more general property like “bounded trackable volume”, which captures perceptual capacities of feature binding and visual tracking (Hatfield 2009).

In short, what an answer to the singular question demands is not only how visual systems manage to segregate and track volumetric structures in a visual array (which is a question for empirical psychology), but also how we as subjects of experience manage to apprehend our external environment as structured into individuals, over and above segregated spatial structures instantiating features. Here’s where the theoretical notion of singular content, as applied to mental states, is particularly useful as a way of marking these capacities. Since perceptualist theories can only give us the former set of capacities, they fail to elucidate the singular question.

According to W.V.O. Quine, what is needed are linguistic abilities. Quine called the process of moving from the experience of a world structured in terms of recurrent instances of properties to the experience of a world structured into enduring physical objects, the ‘reification of bodies’. In a series of writings he described the psychogenesis of reification as a process comprising several stages of increasing cognitive complexity, as the subject masters various linguistic structures like predication, quantification, pronominal reference, relative clauses, etc.⁶

The precise details of Quine’s account need not detain us here; it suffices to say that this account has been severely criticized by philosophers and psychologists

⁶ See Quine 1995 (chapter 3) for a useful summary of the various stages of reification.

alike, due to it being overly intellectualistic and heavily language-dependent.⁷ The main kind of criticism comes from empirical evidence that point to capacities of object individuation in pre-linguistic humans, gathered on the basis of the dishabituation experimental paradigm (Spelke et al. 1995, Carey & Xu 2001, Xu 2007). Since these experiments show psychological abilities of object individuation before the emergence of language, it cannot be that subjects need to master a series of linguistic apparatus if object individuation is to take place.

There is, however, a different problem that may be extracted from Quine's writings. This problem can be formulated in a language-independent manner, and is thus not so easily dismissed by evidence from the dishabituation experimental paradigm. The problem comes from Quine's remarks on a stage of reification where the body is recognized as identical over time across longer stretches of time, which may be seen in the following passage:

"There is still a momentous further step of reification, wherever it may fit in the developmental sequence. It is the transcending of the specious present. Up to that point the reification of bodies is still sketchy, weak in the time dimension. There is as yet no sense in saying that this raven is the one we saw yesterday, or that it is not. We are still dealing with a stage (...) that is limited to the specious present and to short-term memories and expectations." (Quine 1995: 36)

It is at this point that philosophers usually bring in capacities to represent objects within a more comprehensive, objective spatiotemporal framework. As Quine continues:

"The last stage [of reification] is where the body is recognized as identical over time, despite long absences and interim modifications. Such reification presupposes an elaborate schematism of space,

⁷ For representative critiques see Campbell (2002), Bermudez (2003), Carey (2009) and Burge (2010).

time, and conjectural hidden careers or trajectories on the part of causally interacting bodies.” (Quine 1992: 7)

This strategy has been endorsed by P.F. Strawson (1959), Gareth Evans (1982) and more recently by Gary Hatfield (2009). The basic suggestion is that once the subject develops a mature and elaborate scheme of space and time, she acquires the means to maintain, or recover, the object’s numerical identity diachronically through long periods of non-observation. I can tell, for example, that the pocket watch I recently saw in my grandmother’s attic is the same pocket watch I saw my grandfather wearing decades ago because I can trace a plausible spatiotemporal path from p (my grandfather’s pocket) at an earlier time t to p' (my grandmother’s attic) at a much later time t' . The pocket watch’s numerical identity is thus maintained (or recovered) diachronically, thanks to my capacity to speculate about its possible movements and behaviors from place p and time t to a different place p' and a posterior time t' , across all intermediaries times and places in between.

One problem with this solution, however, is that what such a scheme consists in is something notoriously difficult to make precise.⁸ John Campbell, for example, has suggested that our understanding of the simultaneous connectedness of space must be based on our grasp of the movements of objects throughout this space, which of course presuppose a prior grasp of the particularity of objects (1999). Hence, it cannot be that we need a conception of space in order to understand the particularity of objects; the order of explanation would be reversed. Be as it may, a more detailed discussion of this topic is beyond the scope of this paper. But on the basis of these serious difficulties, if we could come up with an answer to the singular question that

⁸ See Bennett (1996) and the essays collected in Eilan, McCarthy and Brewer (1999).

did not require a prior elucidation of an “elaborate schematism of space and time”, that would be preferable.

In the final sections of the paper we will propose a different way of answering this question, which will not require our acquisition of a mature and elaborate scheme of space and time. Although we agree that possession of such a scheme (whatever it may mean) greatly improves our capacity to identify and re-identify objects across successive perceptual encounters, we don't believe this is the only measure of a subject's capacity to apprehend her external environment as structured into particular objects.

We will propose another measure, based on a capacity to actively sustain an object representation in mind across different contexts. This capacity, supported by neural structures in the prefrontal cortex (PFC), allows us to elucidate the singular question in a way that is considerably less intellectualistic than the suggestions sketched in this section. But before we go into details, we will examine two other theories that were also proposed in the literature as middle ground between intellectualism and perceptualism.

IV – Other intermediate theories of singular demonstrative content

To be fair, we are not the only ones to realize that the singular question might demand more than low-level perceptual processes of object segregation and tracking, but less than what intellectualism requires. Imogen Dickie (2010) and José-Luiz Bermudez (2003/2007), for example, have both proposed intermediate theories between these two extremes, which will be briefly examined in turn.

Although Dickie happily accepts the perceptualist claim that low-level perceptual processes are able to correctly select objects in a visual array, in order for us to be able to single them out in thought other abilities are required. Her motivation for postulating these extra set of abilities is to differentiate between actual material objects and things like shadows and water ripples. For at the level of the visual system, both are considered as ‘visual objects’, i.e., items that are segregated and processed by the visual system as units. So if at the level of the visual system there is no differentiation between material objects and shadows or water ripples, how can we refer to objects based on this perceptual capacity alone?

Her solution is to say that perceptual abilities to single out and track perceptual units must be supplemented by a subject’s sensitivity to the possible ways in which objects might behave, according to their category (i.e, material object). In other words, when you attribute a property to some object – say, the property of being solid – this attribution is governed by a certain ‘template’ you have of that object, which determines the kind of properties it is possible for it to have, appropriate to its category.

Templates are dynamic notions that only make sense diachronically, in terms of the cluster of dispositions you have concerning the object’s states and doings over time. When you attribute the property of solidity to an object of perception, we cannot know only on the basis of this single property attribution whether you have managed to single out an object in thought or not. For suppose at time t you attribute the property of solidity to it, but at time t' the object does something completely unexpected such as vanish out of thin air. When faced with this violation of solidity, you should be disposed to withdraw your earlier attribution, in order to maintain the coherence and stability your template of material objects. Failure to do so reveals that

you have a faulty template, a template that allows for violations of solidity that are not possible for material objects. Hence, you do not have a singular demonstrative thought about it; it is a case of reference failure.

In a similar vein, Bermudez argues that objects are more than just perceptual units that can be segregated and tracked in a visual array by low-level perceptual mechanisms. In his ‘object properties model’ (2003/2007), to perceive an object is to have a perceptual sensitivity to what he calls ‘canonical object properties’. These properties circumscribe what it is for something to count as a physical object, and correspond to certain physical principles and regularities that govern the behavior of objects. For example: objects tend to be resistant to pressure, to be subject to the effects to gravity, to have a certain mass that govern their interactions with other objects, and so on (2003:82). To each of these regularities, there corresponds a higher-order property these objects possess: for example, the property of being resistant to pressure, of being subject to gravity, etc. Unless one is perceptually sensitive to these properties, manifested in certain expectations and anticipations regarding the object’s behavior as one actively interacts with it, one would not be able to single it out in perception or thought.

Like Dickie, Bermudez complements low-level perceptual processes with certain aspects of subjects’ dispositions and active interactions with objects in the world. In this sense, both may be understood as theories of singular demonstrative content that are intermediate between perceptualism and intellectualism.

We have sympathies for these accounts – specially if we interpret Dickie as proposing a less rationalist and more inclusive notion of property attribution, in dispositional terms for example – where to attribute the property ‘...is edible’ to an

object is nothing more than to attempt to eat it. They allow us to take one step beyond standard perceptualist theories that mention only low-level perceptual processes, and bring into the picture important elements of subjects' actions and dispositions towards material objects. We believe this is plausible and on the right track, and our account should be understood more as a way of complementing these theories than as philosophical competitors.

But because they are still restricted to contexts of perceptual interaction, they fail to fully answer the intellectualist challenge, which brings out the worry that singular demonstrative content stretches beyond the context of perception. Our account allows us to take one step even further, going beyond perception while still falling short of more sophisticated intellectual abilities. The details of this account will be fleshed out in section V.

V – Singular demonstrative content and the prefrontal cortex

The prefrontal cortex (PFC) is a collection of interconnected neocortical areas that sends and receives projections from virtually all cortical sensory systems, motor systems and subcortical structures (Miller & Cohen 2001). The PFC has been linked to analogical reasoning, working memory, cognitive control, decision-making and problem solving (Speed 2010, O'Reilly & Frank 2006). Kharitonova and Munakata have shown for example that perseverance behavior in task-switching (i.e., failure to modify one's past behavior in light of the novel task) is linked to task-relevant information being represented in posterior cortical areas, which has been

independently shown to encode information in a more detailed and stimulus-specific manner.⁹

More flexible switching and adaptation, in contrast, is linked to task-relevant information being actively sustained *via* neuronal firings in the prefrontal cortex, allowing it to be maintained in working memory as one switches flexibly from one task to another. This is because prefrontal representations encode information in a more abstract format, collapsing across specific idiosyncratic details and forming more abstract superordinate categories, which allow them to be more easily applied outside the learning context.

But what exactly is the role of the prefrontal cortex in creating these more abstract superordinate categories? We can make this idea clearer with Nicolas Rougier (and collaborators') work on connectionist models of PFC processing (2005). Their hypothesis is that activation of PFC-specific neurobiological mechanisms in a wide range of task experiences is sufficient for the PFC to self-organize into more abstract representations, allowing the network to make more successful generalizations to novel circumstances.

The stimuli presented to the network varied along a number of dimensions (shape, size, color and spatial location), and the networks were tested in either a subset or the totality of tasks directed at naming or comparing different features. The authors tested networks of varying complexity and under different conditions: with all or a subset of the PFC-specific neural mechanisms, and with training in two or all of the four possible tasks.

⁹ See Kharitonova & Munakata 2011, Miller & Desimone 1994, Jog et al. 1999

Interestingly, what they found is that only the full PFC model, encompassing all PFC-specific mechanisms and capacities, developed synaptic weights and patterns of activity that encoded more abstract representations of dimensions of features ('color', 'shape', etc.), which encompass various features within any given dimension (Rougier et al. 2005: 7340). These representations aided task performance by being actively maintained in the PFC while sending top-down excitatory support for the relevant stimulus dimension: for example, if the task at hand was to name a particular color, the more abstract representation of the 'color' dimension would be actively maintained in the PFC, and would constantly send excitatory signals to sensory representations of particular colors for the whole duration of the task. An adaptive gating mechanism, in turn, learned to update PFC activity whenever the relevant dimension to the task at hand switched. If the task changes from naming colors to naming shapes, a more abstract representation of the 'shape' dimension would be actively maintained, and excitatory signals sent to the sensory and posterior cortices.

Due to the more abstract nature of these representations, we should also expect that they would be able to modulate not only processing of familiar stimuli (i.e., those in which the network was trained) but also processing of novel stimuli – that is to say, we should expect the full PFC network to be able to generalize from its prior experience to novel situations, which is exactly what was observed. The full PFC network exhibited significant generalization capacities, achieving 85% accuracy on stimulus for which it had no prior same-task experience (Rougier et al. 2005: 7341).

In addition, significant generalization was observed only in networks that received training in all the tasks: those that trained only in a task subset showed reduced generalization capacities, making as much as four times more errors. This shows that cognitive flexibility and generalization arise not only from a fully

operational PFC but also as a function of the range of experiences we have in the world. As we actively engage with our external environment in a variety of tasks, we bring to the foreground of our experience now one dimension of the environment, now another; on this basis, PFC-specific neurobiological mechanisms are able to stabilize together certain features of the environment into more abstract superordinate dimensions, which endows the cognitive system with greater cognitive flexibility and adaptability to novel situations.

And what we now want to suggest is that ‘object’ is precisely one of these categories: an abstract, superordinate PFC representation that emerges as we actively interact with objects in the world and learn of their various behaviors. This representation emerges gradually in the PFC, as a result of a wide range of experiences with objects in the world combined with PFC-specific neurobiological structures.

We have briefly seen in section IV Bermudez’s ‘object properties model’, where objects are understood as things that behave according to certain physical regularities. We agree. As we actively engage with objects in various ways we bring some of these properties to the forefront of our experience, depending on the task at hand: that objects are subject to the effects of gravity, that they have a certain mass which govern their interactions with other objects, etc. And just like the network in Rougier’s experiments has learned to extract a more abstract dimension of ‘color’ from a series of interactions with particular colors, the more abstract category ‘object’ is extracted from repeated patterns of interactions with objects, as we experience their difference properties. For example, if one is engaged in a task that highlights the property of objects of being subject to gravity, and then shifts to a context where the property of objects of causally influencing other bodies is highlighted, a more abstract

representation of 'object' will be activated in the PFC, sending excitatory signals to the posterior and sensory cortices and guiding our active interactions with the object in each case.

The more abstract nature of these representations also allows us to generalize from our prior experience to novel encounters, as the PFC representation activates whenever we are faced with (what seems to behave like) an object. This representation may be actively held in mind even if the object cannot be currently perceived, allowing us to represent the continued existence of the object across perceptual and non-perceptual contexts.

Instances of this capacity have been observed in several delayed matching to sample tasks, where subjects need to actively sustain an object representation in mind in order to match it to a target object that appears after a certain time, during which the object is unperceived (Miller et al. 1996, Warden & Miller 2007). What these studies show is that cells in the PFC that represent the target object show heightened activity during the delay period, which is taken to be the neural correlate of the capacity to represent the continued existence of the object during longer periods of non-observation.

At the same time that activity in the PFC is enhanced, activity in sensorimotor areas that represent more specific surface properties of the object is diminished during the delay period (Courtney et al. 1998). Hence the emphasis we put on PFC-specific object representations as a way of addressing the singular question, over and above sensory representations of objects emphasized by perceptualist theories of singular content. During periods in which the object is not currently perceived the sensory representation of the target object will weaken or disappears altogether, but the more

abstract PFC representation will allow us to represent the object's continued existence beyond the context of perception.

These effects, as Rainer and Miller argue, aren't due to attentional modulation but to what they call 'experience-based modulation'; that is to say, through our experiences with objects selective strengthening occurs between neurons in the PFC and neurons in the visual system that represent the object in question. This allows the representation to be actively maintained through heightened activity in the PFC even in the absence of perceptual input, for increasingly longer periods of non-observation (Rainer & Miller 2000: 186).

VI. Conclusion

We are now finally in a position to give an adequate answer to what we have called the 'singular question' concerning singular demonstrative content. If we imagine a case where a subject is perceptually confronted with an individual object and then goes through a change of context – i.e., the object is removed from sight for a certain period of time – the PFC affords her a way to actively maintain the object representation in mind until she encounters it again, maintaining its numerical identity beyond the context of perception.¹⁰

For example, as we experience a familiar object travel away from us and later return, or as we travel away from the object and return at a later time, we can exploit

¹⁰ Granted, sensory representations of objects may also be said to “move beyond the context of perception”, in the sense that they survive temporary occlusions behind obstacles, and thus can be maintained in the absence of actual perceptual input. But in this case the representation is short-lived (no more than 8 seconds according to research conducted by Noles et al. 2005) and dependent on perceptual expectations concerning the object's reappearance at the other end of the occluder. But current research shows that the more abstract nature of PFC representations allow it to survive longer periods of non-observation in a way that is less dependent on the context of perception.

our capacity to keep the object's representation active in the PFC in order to maintain the object's numerical identity diachronically. A child who has seen its Christmas present be placed under the tree on Christmas Eve can actively maintain its representation in mind throughout the night, as she lies awake anxiously waiting for Christmas morning. When she returns after a sleepless night she presumes she is faced with the same object left under the tree the previous night, which she kept in mind during this interval.

Granted, our capacities for temporal and spatial representation at this level may not be sufficiently developed to support successive re-identification of the object over long stretches of time. But these capacities, we maintain, suffice to distinguish individual objects from instances of property-kinds as the *representata* of our mental states. For representations of the property-kind 'bounded trackable volume' are sensory representations that fade away shortly after the perceptual encounter has come to an end. Flexible PFC representations, in contrast, may be actively maintained for longer periods of time, allowing us to maintain an object's numerical identity beyond the context of perception.

Of course, this capacity is fallible, as it is possible for the object to be replaced by a qualitatively identical one at some point during periods of non-observation. In this case, our re-identification of the object as the same one we have perceived earlier (and which we have kept in mind) will turn out to be mistaken. But there's nothing wrong with this scenario: developing further ways to check for sameness of numerical identity after a period of interrupted observation is precisely one of the ways in which we expect our knowledge of objects to develop. Our ability to draw a distinction between numerical and qualitative identity is itself a gradual matter, which develops in time as we continue to interact with objects in the world in various ways. There

needn't be a magical point where maintenance of numerical identity across different contexts suddenly becomes possible.

To conclude, we have started out with two questions, which any putative theory of singular demonstrative content must be able to answer:

- **The demonstrative question:** what makes the contents of these mental states 'demonstrative', i.e., connected to things in the external world?
- **The singular question:** what makes these contents 'singular', i.e., contents that represents individual objects as opposed to properties, events or some other suitable candidates for the *representata* of these mental states?

Regarding the first, we have granted that perceptualist theories have adequate resources to address the demonstrative question in a satisfactory way. As we know more about how our visual system is able to extract and process visual information, we will be able to answer the question of how the mind, via visual perception, can make contact with external things in the world.

But we have also said that the perceptualist strategy has its limitations when it comes to the second, singular question. Individual objects are more than spatially segregated volumes that may be visually tracked during a period of observation. They are also enduring physical things that persist in time, and that continue to exist after the perceptual encounter has come to an end. Perceptualist theories, which attempt to answer the singular question only on the basis of sensory representations of objects, face limitations when confronted with this problem.

One answer to the singular question, which may be found in Quine, Strawson, Evans and Hatfield, emphasizes our acquisition of a mature and elaborated conception

of space and time, required to speculate about the possible movement and behavior of objects during long periods of absence.

We do not doubt that such a conception of space and time does indeed offer a solution to the singular question. But we dispute that this is the only possible solution. Another option, which is considerably less intellectualistic, lies in a capacity to actively sustain an object representation in mind through longer stretches of time even in the absence of perceptual input. This capacity is supported by more abstract representations that are formed in the PFC. These representations emerge gradually and naturally as PFC-specific neurobiological structures are exposed to a wide range of experiences with objects in the world.

This capacity allows us to differentiate between individual objects and instances of the property-kind 'bounded trackable volume' as the *representata* of our mental states, without having to appeal to more sophisticated forms of spatiotemporal representations. And when we become able to represent individual objects in this manner, we are then in a position to engage in a series of object-directed mental activities with singular demonstrative contents. The singular element of this content is accounted for by our capacity to actively sustain an object representation in mind while detached from the context of perception, whose neural correlate is heightened activity in cells in the PFC that represent the target object.

But before we bring our discussion to an end, we must admit that perhaps both sides would be unmoved by our account. That is to say, perhaps philosophers of a perceptualist inclination will insist that there is no need to go beyond the context of perception (although there may be good reasons to go beyond low-level sensory processes), while intellectualists will protest that this proposal is too weak; that the

notion of object persistence that is needed here is even stronger than the one we rely on, as entities that may be re-identified over very long stretches of time and across different places. This is something that might indeed require more sophisticated spatiotemporal abilities.

But there is a way of conciliating both worries. As we see it, capacities to single out particular objects in thought might come in degrees, and there may be no specific set of abilities that demarcate a precise point where singular demonstrative thought suddenly becomes possible. We can then see all the abilities that were discussed throughout the paper as indications that one is getting hold of the capacity to single out particular objects, and that one is developing one's understanding that the world is structured into individual things. In this picture, the perceptualist would be correctly describing the lower end of this spectrum of abilities, while intellectualists would be describing its higher end. Our account, in turn, would be something in between.

So, to be clear, we are not claiming that singular demonstrative thought is to be analyzed exclusively at the level of the representational abilities associated with the prefrontal cortex. Our aim is more modest than that. We are rather bringing these abilities into the discussion and pleading for their incorporation into the theory of singular demonstrative content, while at the same time drawing attention to a body of empirical research that has been largely ignored by both perceptualists and intellectualists. We believe this research may provide us with important insights concerning singular demonstrative representation. With this paper, we hope to have provided the means to take singular demonstrative thought one step further beyond contexts of perceptual interaction, while still falling short of more sophisticated intellectual abilities. In this sense, our account may be of interest to both perceptualists and intellectualists alike.

References:

- Bennett, J. (1996). Do Animals Have Cognitive Maps? *The Journal of Experimental Biology* 199 (Pt. 1), 219-224.
- Bermudez, J. L. (2003). *Thinking Without Words*. Oxford: Oxford University Press.
- Bermudez, J. L. (2007). The Object Properties Model of Object Perception: Between the Binding Model and the Theoretical Model. *Journal of Consciousness Studies*, 14 (s 9-10), 43–65.
- Brewer, B. (1999). *Perception and Reason*. Oxford: Oxford University Press.
- Burge, T. (1977). Belief *De Re*. *Journal of Philosophy*, 74 (6), 338-362.
- Burge, T. (2010). *The Origins of Objectivity*. Oxford: Oxford University Press.
- Campbell, J. (1999). The Role of Physical Objects in Spatial Thinking. In: Eilan, N., McCarthy, R., & Brewer, B. *Spatial Representation: Problems in Philosophy & Psychology*. Oxford: Oxford University Press.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.
- Carey, S., & Xu, F. (2001). Infants' Knowledge of Objects: beyond Object Files and Object Tracking. *Cognition*, 80 (1-2), 179-213.
- Courtney, S.M., Petit, L., Maisog, J. M., Ungerleider, L.G., & Haxby, J. V. (1998). An area specialized for spatial working memory in human frontal cortex. *Science*, 279, 1347-1351.

- Cussins, A. (1992). Content, embodiment and objectivity: the theory of cognitive trails. *Mind* 101 (404). 651-688.
- Dickie, I. (2010). We are acquainted with ordinary things. In: Jeshion, R. *New Essays on Singular Thought*. Oxford: Oxford University Press. 213-245.
- Eilan, N., McCarthy, R., & Brewer, B. *Spatial Representation: Problems in Philosophy & Psychology*. Oxford: Oxford University Press.
- Evans, G. (1983). *The Varieties of Reference*. Oxford: Oxford University Press.
- Hatfield, G. (2009). *Perception and Cognition: Essays in the Philosophy of Psychology*. Oxford: Oxford University Press.
- Jog, M. S., Kubota, Y., Connolly, C. I., Hillegaart, V., & Graybiel, A. M. (1999). Building Neural Representations of Habits. *Science* 286 (5445), 1745–1749.
- Kharitonova, M., & Munakata, Y. (2011). The Role of Representations in Executive Function: Investigating a Developmental Link Between Flexibility and Abstraction. *Frontiers in Psychology*, 2, 347.
- Levine, J. (2010). Demonstrative Thought. *Mind and Language*, 25 (2), 169-195.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. *The Journal of Neuroscience*, 16 (16), 5154 –5167.
- Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, 24, 167–202.
- Miller, E. K., & Desimone R. (1994). Parallel Neuronal Mechanisms for Short-term Memory. *Science*, 263 (5146), 520–522.

- Noles, N. S., Scholl, B. J., Mitroff, S.R. (2005). The Persistence of Object File Representations. *Perception & Psychophysics* 67 (2) (February), 324–334.
2005.
- O'Reilly, R.C., & Frank, M. J. (2006). Making Working Memory Work: A Computational Model of Learning in the Prefrontal Cortex and Basal Ganglia. *Neural Comput.* 18 (2), 283–328.
- Pylyshyn, Z. (2001). Visual Indexes, Preconceptual Objects, and Situated Vision. *Cognition*, 80, 127–158.
- Pylyshyn, Z. (2007). *Things and Places: How the Mind Connects With the World*. Cambridge, MA: MIT Press.
- Quine, W.V.O. (1992). Structure and Nature. *Journal of Philosophy*, 89 (1), 5-9.
- Quine, W.V.O. (1995). *From Stimulus to Science*. Cambridge, MA: Harvard University Press.
- Rainer, G., & Miller, E. K. (2000). Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, 27 (1), 179-189.
- Rougier, N. P., Noelle, D. C., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (2005). Prefrontal Cortex and Flexible Cognitive Control: Rules Without Symbols. *Proceedings of the National Academy of Sciences of the United States of America*, 102 (20), 7338–7343.
- Speed, A. (2010). Abstract Relational Categories, Graded Persistence, and Prefrontal Cortical Representation. *Cognitive Neuroscience*, 1 (2), 126–137.
- Spelke, E.S., Kestenbaum, R., Simons, D.J., & Wein, D. (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, 13 (2), 113-142.

- Strawson, P.F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- Treisman, A. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology* 12 (1). 97-136.
- Warden, M. R., & Miller, E. K. (2007). The Representation of Multiple Objects in Prefrontal Neuronal Delay Activity. *Cerebral Cortex*, 17, 41-50.
- Wu, W. (2011). What Is Conscious Attention? Philosophy and Phenomenological Research 82 (1): 93-120.
- Xu, F. (2007). From Lot's Wife to a Pillar of Salt: Evidence that 'Physical Object' is a Sortal Concept. *Mind & Language*, 12 (3-4), 365-392).