

Object Files, Properties, and Perceptual Content

Santiago Echeverri¹

Published online: 12 June 2015

© Springer Science+Business Media Dordrecht 2015

Abstract Object files are mental representations that enable perceptual systems to keep track of objects as numerically the same. How is their reference fixed? A prominent approach, championed by Zenon Pylyshyn and John Campbell, makes room for a non-satisfactional use of properties to fix reference. This maneuver has enabled them to reconcile a singularist view of reference with the intuition that properties must play a role in reference fixing. This paper examines Campbell's influential defense of this strategy. After criticizing it, a new approach is sketched. The alternative view introduces representational contents to explain perceptual individuation. After arguing that those contents are not satisfactional, it is concluded that there is room for a third view of reference fixing that does not fit into the singularist/descriptivist dichotomy.

There are two traditional ways of explaining the determination of reference. According to descriptivism, reference is fixed by an object's satisfaction of properties. On the singularist view, it is fixed by causal relations to objects (Bach 1987; Burge 1977; Recanati 2012).

Early work on reference includes a number of attempts at reconciling these two views. A familiar strategy consists in introducing causal elements into the descriptions that fix reference (see Recanati 2012 for discussion). More recently, a new line of research has emerged in the psychology of reference. Focusing on perception, some authors have tried to reconcile the singularist view with the widespread intuition that properties must play a role in the determination of reference (Campbell 2002, 2011, 2012, 2013; Pylyshyn 2007). This has led them to introduce non-satisfactional uses of properties.

This paper has two goals: It examines Campbell's singularist account and proposes a new approach that posits non-satisfactional contents to fix reference. To this end, I submit that non-representational views face circularity problems that can be avoided by positing reference-fixing contents. Since these contents refer in a non-satisfactional

✉ Santiago Echeverri
santiagoecheverri@hotmail.com; <http://www.unige.ch/lettres/philo/episteme/>

¹ Faculté des Lettres, Département de Philosophie, Université de Genève, 2, rue de Candolle, CH-1211 Genève 4, Switzerland

manner, they offer an account of reference fixing that goes beyond the descriptivism/singularism dichotomy.

The paper falls into seven sections. Section 1 introduces the concepts of a mental file and reference fixing. Section 2 frames the opposition between descriptivism and singularism. Section 3 explains how one could make room for a non-satisfactorial use of properties in the fixation of reference. In Section 4, I submit that it is reasonable to frame the non-satisfactorial use of properties by positing *sui generis* representational contents that underlie perceptual individuation. In Sections 5 to 7, I make some constructive remarks on the structure of the contents proper to perceptual individuation and compare the resulting view with traditional proposals.

1 Mental Files and Reference Fixing

I will frame the discussion with the notion of a mental file, which has played an influential role in theorizing on reference (Dickie 2010; Fodor 2008; Recanati 2012). A mental file is a mental representation analogous to a singular term that enables a system to store information about individuals. Although mental file theorists often disagree on the nature of files, what I will say here is compatible with most accounts. I will add only the distinction between ‘object files’ and ‘recognitional files,’ which has been widely used in contemporary work on perceptual reference. According to most cognitive scientists working on this topic, there is a sharp difference between keeping track of an entity as numerically the same and recognizing it. On their view, these two tasks are underwritten by ‘object files’ and ‘recognitional files’ respectively (Kahneman et al. 1992: 176–7).

Imagine that someone is approaching you from a long distance. At some point, you are able to identify that person as an old friend. What happened during the whole episode? You visually followed a portion of reality as numerically the same. After a while, you were able to *recognize* that moving body as your friend. We can therefore say that, in some cases, perceptual tracking and recognition come apart. The mental file theorist says that, during the whole episode, an object file was activated to represent the moving body as numerically the same. That object file is a representation in working memory that became active to store perceptual information about that moving object. When you recognized the moving object as your friend, the object file was linked to another, recognitional file, let us say PIERRE. This file might contain long-term memories about Pierre’s face, his age, his name, and so on.¹

In this paper, I shall examine the conditions under which the reference of object files is fixed. In the philosophy of language, what fixes the reference of an expression *e* is the set of facts that make it the case that *e* refers to a specific entity and not to anything else (Kripke 1980).² This question can be easily generalized to mental files if one thinks that they are similar to singular terms that purport to refer to individuals. We can

¹ I shall use capitals to denote mental representations, single quotes to denote words, and angle brackets ‘⟨...⟩’ to denote properties or representational contents.

² I am using the word ‘fact’ in a broad sense, to include any factor that may be relevant to explain why a representation has the referent it presently has.

therefore ask: What sorts of facts make it the case that a mental file M refers to an entity e and not to anything else?

A number of cognitive scientists think that humans (and other animals) have innate mechanisms of object perception (Carey 2009; Pylyshyn 2007; Spelke 1990). Thus, humans (and other animals) have an innate repertoire of object files they can activate in working memory to store information about objects. Current models suggest that infants can activate up to three object files simultaneously, while adults can activate up to four or five. If something along these lines is correct, the reference-fixing question for object files is: What sorts of facts make it the case that an object file O refers to an object o and not to anything else?

2 The Descriptivism/Singularism Dichotomy

In his recent work on mental files, Recanati (2012) offers a reconstruction of the history of the theory of reference as the opposition of two camps: descriptivism and singularism. These views can be understood as providing different accounts of reference fixing. In what follows, I spell out their strategies to account for reference fixing and explain how they apply to perception.³

How can we cash out the descriptivism/singularism dichotomy in the case of perception? A simple way of drawing the line is to present it as a question of priority. According to the descriptivist, “our [perceptual] relation to individual objects goes through properties of those objects. What are given to us are, first and foremost, properties whose worldly instantiation we are able to detect, and only indirectly objects.” (Recanati 2012: 3) According to the singularist, the order of priority goes in the other direction: “In perception, we are related to the object we perceive. The perceptual relation is what enables us to gain (perceptual) information from the object.” (Recanati 2012: 20; see also Bach 1987: 12)

This opposition suggests two models of reference fixing. Descriptivists defend a properties-first model (PFM) according to which reference must go through properties. Singularists advocate an acquaintance-first model (AFM) that depicts reference as a relation to objects.⁴

I will take for granted a realist view of properties. Since nothing of what I will say depends on whether properties are universals or tropes, I will remain neutral on their nature. In order to avoid some possible pitfalls, I will restrict the discussion to relatively uncontroversial properties like color or shape. Our question concerns the way properties play their *mediating role* in descriptivism. When Recanati formulates the opposition, he has in mind a specific model of property mediation. PFMs endorse the conjunction of two claims:

³ This dichotomy shapes influential views like those of Bach (1987: 12) and Burge (1977: 346), among many others.

⁴ This debate is orthogonal to the question of whether we are also acquainted with properties. One can term ‘pluralist’ any view according to which we are both perceptually related to objects *and* to properties. Many acquaintance theorists are pluralists, for they posit perceptual relations to colors, shapes or size (Campbell 2007; Kennedy 2007; Russell 1912).

Representational claim: The representation of the properties of objects is a pre-condition for referring to those objects.

Satisfactional claim: The determination of reference follows the model of Russell's (1905) theory of definite descriptions, which imposes a uniqueness condition on the instantiation of properties. If a property F fixes the reference of an object file O :

- (1) There must be *at least* one object that is F .
- (2) There must be *at most* one object that is F .

When we conjoin these two claims, we get a model of reference fixing for object files. If an object file contains some predicates $P_1, P_2, \dots, P_{n-1}, P_n$, it will refer to whatever entity satisfies $P_1, P_2, \dots, P_{n-1}, P_n$, most of them, or a weighted sum of them. Recanati quotes Chastain's (1975: 254) apt analogy: "we get at physical objects only by a semantic shot in the dark: we specify properties or relations and hope that they are uniquely exemplified."⁵

In what follows, I refer to accounts of reference fixing that endorse the two claims as 'satisfactional.' So 'non-satisfactional' will be used to characterize those views that reject either of the two claims.⁶

AFMs reject the two claims by introducing the notion of acquaintance with objects. In a *weak sense*, acquaintance denotes *any* way of gaining knowledge of things that is not satisfactional. Let us call this the 'weak acquaintance-first model' (WAFM). In a *strong sense*, it denotes a *specific* alternative to satisfactional reference fixing. Hereafter, I will call it 'strong acquaintance-first model' (SAFM).

One way of developing a SAFM would be to claim that properties play no role in the fixation of reference. This view does not seem promising. After all, one might argue that properties must play, at the very least, a *causal* role in the fixation of reference. If one perceives an object o , one must be standing in a causal relation to o . Thus, standing in a causal relation to o is necessary in order to perceive o . On a standard analysis, causation is a relation between events or states of affairs. On this view, it is the objects *as possessors of some properties* that produce changes in the world (Shoemaker 1980). Hence, if perceptual relations involve causation, one is committed to making room for properties in the fixation of reference.

The singularist could take advantage of this line of thought to provide an amended version of SAFM. Recall the claim that the representation of the properties of objects is a pre-condition for referring to those objects. The singularist might point out that, in order to influence our perceptual systems, the properties of objects need not be represented (Pylyshyn 2007). On most views, perceptual acquaintance requires a causal relation to objects (Campbell 2002; Recanati 2012). Thus, the singularist only needs to introduce properties into the characterization of those causal relations. Hence, the

⁵ A slightly different view analyzes perceptual content as having the existential form $\exists x(Fx)$. Defenders of this view drop the uniqueness condition (Davies 1992; Jackson 2012; McGinn 1982). I examine this approach in Section 7.

⁶ Some philosophers use the phrase 'satisfaction conditions' as synonymous with representational content (Searle 1983). This is not the intended meaning. In Sections 4–6, I argue that there are non-satisfactional representational contents, i.e., contents that do not fix reference through a mechanism that complies with the representational and satisfactional claims.

priority of acquaintance over properties does not mean that the latter play no role in explaining the fixation of reference. It simply means that, *pace* descriptivism, the properties of objects do not mediate reference by being represented.

In what follows, I examine Campbell's version of SAFM. Campbell combines a rejection of the representational claim with the contention that properties do play a causal role in perceptual reference. I will suggest that Campbell's model does not offer an adequate account of reference fixing (Section 4). Subsequently, I shall use the results of the discussion to formulate an alternative view (Sections 5–7).

3 Campbell's Selection/Access Model

Campbell's view goes beyond the causal model sketched in Section 2 because he thinks that *conscious attention* also plays a central role in reference fixing. By this he means two things: (1) that knowledge of reference requires consciousness and (2) that properties enable us to select objects. I propose to bracket the vexed issues of consciousness and attention in order to explore the role of properties in selection.⁷

By focusing on selection, Campbell seeks to satisfy two conditions: make room for properties in the individuation of objects and avoid the appeal to representational contents in the fixation of reference. In *Consciousness and Reference*, he distinguishes between:

- (1) Using an object's property to single it out visually, and
- (2) Verifying a proposition to the effect that the object has that property (Campbell 2002: 29; see also: 30, 33).

In the same context, he insists that (1) is more primitive than (2), for one can use a property to single out an object without having a concept of that property. He makes this point by invoking the cases of animals and pre-linguistic infants that lack concepts of color but use color vision to segregate objects from their background (2002: 29).

This distinction provides a welcome refinement of the causal model. By making room for the use of properties to single out objects, Campbell specifies how properties can contribute to the activation of object files. To see how the account works, consider the Ishihara Color Vision Test (Campbell 2012: 110-ff.; 2013: 822-ff.). In this test, a subject is presented with various dots of different color and size. The test works because some of the dots form a pattern like a numeral. This allows for identifying color deficiencies if the subject cannot detect the pattern or has difficulties in seeing it. The test thereby conjoins grouping principles and color perception.

If you are a normal trichromat, you will see a green '5' in Fig. 1. If you are not a normal trichromat, however, you will be perceptually acquainted only with the dots on the plate. Thus, your visual system will be unable to activate an object file for the green '5.' Clearly, in order to individuate the '5,' one must be sensitive to the contrast in hue between some dots and the others. Thus, sensitivity to some color properties is *necessary* to individuate the '5.'

⁷ It has been argued that there is pre-conscious and/or pre-attentive selection of objects. See Pylyshyn (2007).

This case illustrates how a relatively uncontroversial property like color can be used to single out an item in perception. In more recent papers, Campbell (2011, 2012, 2013) revisits this distinction by appealing to a recent empirical model of attention by Huang and Pashler (2007). On Campbell's view, there is a fundamental contrast between "selecting a region or object by using some property of it, such as its color or texture" and "accessing a property of that selected region or object." (2011: 324)

Campbell follows an important tradition—probably initiated by Block (1995)—according to which access is a functional notion. Thus, a property F is accessed by a system S just in case S uses F . F is accessible to S just in case F is available for use by S . And F is available for use by S just in case S can perform inductive generalizations concerning F , can report the instantiation of F on the basis of perception, can form a belief concerning F , etc. (Campbell 2002: 29; 2007: 11–2; 2011: 331; 2012: 111; 2013: 820–ff.). Consider some examples. We can use the color of a salad to predict whether it is fresh and a doctor can use a patient's skin coloration to determine whether she is ill (Campbell 2007: 12). By contrast, one could use the color of an object to segregate it from its background but form no beliefs about the relevant color (Campbell 2013: 822–3).

On Campbell's view, none of these further uses is conceptually necessary for the use of color to individuate objects. Whereas one can offer an exhaustive characterization of access to a property in functional terms, the only way of characterizing the use of a property for selection is to cite its role in individuation (Campbell 2011: 332; 2013: 823).

This separation from further uses is crucial to Campbell's project because it enables him to block Siegel's (2010) influential argument in favor of perceptual contents. That argument starts from the premise that properties are presented in perception and concludes that perceptual experience has representational content. If one can use properties only for selection, however, one can avoid Siegel's content view: "The accessing of a property by the perceiver may indeed be a matter of representational content; but the figuring of a property in visual experience is more basic than its being accessed by the perceiver." (Campbell 2013: 820)

I want to remain neutral on whether accessibility is a necessary condition for representational content. Nevertheless, I will show that there are good independent reasons to characterize perceptual individuation by positing representational contents. If my view is correct, there are non-satisfactional representational contents that *explain* the fixation of the reference of object files. Thus, my picture provides the materials for a conception of reference fixing that does not fit into the singularism/descriptivism dichotomy.

4 Perceptual Individuation Has Representational Contents

In this section, I argue that perceptual individuation is best explained by positing representational contents. Before I present the argument, some preliminary remarks are necessary.

I will use the noun 'representation' to denote a representational vehicle. Thus, the string 'John is a dog'—phonologically or morphologically individuated—is a representation. By analogy, its psychological counterpart JOHN IS A DOG is a *mental* representation. I will say that a mental representation *carries* a representational content.

Following a long tradition, a representational content will be understood as an abstract entity. Typical examples of contents are propositions and their constituents. As will become clear later, these entities do not exhaust the class of representational contents. I shall say that representational contents determine *correctness conditions*. The latter are situations under which a representational content is correct or incorrect. The proposition ⟨John, being a dog⟩ is correct if John exemplifies the property of being a dog. It is incorrect otherwise. I will use the adjective ‘correct’ as a generic term covering a variety of evaluations like truth or accuracy.

Some singularists assert that perceptual states are not the sorts of things that can be correct or incorrect because correctness conditions only arise at the level of propositional attitudes like belief (Brewer 2011). A less radical form of singularism grants that it is appropriate to classify perceptual states as correct or incorrect. Nevertheless, this association of correctness conditions with perceptual states lacks explanatory bite (see Pautz 2011 and Schellenberg 2014 for discussion). Thus, singularists can be reluctant to analyze perceptual states in terms of representations or representational contents while still associating correctness conditions with perceptual states. Hence, in order to make a substantial claim against this brand of singularism, one has to analyze perceptual states as relations to representational contents.

Campbell’s view fits pretty well with the less radical attitude. He could grant that the use of properties for selection may be associated with correctness conditions. Yet, he thinks that relations to objects are more fundamental than representational contents in the characterization of perceptual states. So he would be reluctant to analyze perceptual states as relations to representational contents.

One may wonder what it would take to make a non-trivial use of representational contents in the explanation of reference fixing. Consider an analogy from Schellenberg (2014). One can describe a painting as representing that things are thus and so. Yet, it does not follow that the painting has the content expressed by one’s description. Although one’s description can be informative, this will not tell us anything substantial about the way that painting refers to the world. Thus, if reference fixing is underwritten by representational contents, those contents are not only useful descriptions of perceptual reference but also contribute to explaining how object files acquire their referents.

Our question is: What sorts of facts make it the case that an object file *O* refers to an object *o*? The acquaintance theorist responds: If a subject stands in a perceptual relation to *o*, the resulting object file *O* refers to *o*. Campbell offers a more sophisticated picture: It is not only a relation to an object that fixes reference. After all, properties also play a causal role in the selection of the object. Yet, Campbell avoids the introduction of representational contents to characterize selection because the latter only intervene *downstream* of perception, i.e., when one makes a further use of the selected property.

In what follows, I want to make a case for a different story. In the visual modality, what fixes reference is the fact that the subject correctly represents a multiplicity of elements as *one volumetric* entity. This representational content must be located *upstream* of the activation of object files. As far as vision is concerned, the activation of some primitive representations is a pre-condition for standing in a perceptual relation to an object.

I will present a three-step argument to defend this picture. First, I will introduce the concept of perceptual individuation and show that it is underwritten by at least two different sorts of perceptual organization: grouping and figure-ground segregation.

Second, I will show that it is natural to associate correctness conditions with these two sorts of perceptual organization. Interestingly, these correctness conditions differ from those often discussed in the philosophy of perception, which typically involve the attribution of properties to objects. Third, I will suggest that this characterization of perceptual organization in terms of correctness conditions is not only descriptively appropriate but also explanatorily relevant. If my considerations are plausible, perceptual organization is not only associated with correctness conditions but is also explained by representational contents that determine those correctness conditions.

4.1 Individuation and Perceptual Organization

Campbell's emphasis on selection is reminiscent of psychologists' concept of individuation (Carey 2009; Pylyshyn 2007). Individuation is a process that enables perceptual systems to 'detach' various elements from the scene as one object. According to most views, individuation requires perceptual organization. A widespread hypothesis in cognitive science is that there are various kinds of perceptual organization that occur at different processing levels (Peterson and Kimchi 2013; Wagemans *et al.* 2012). Two forms of perceptual organization are relevant to the following argument: *grouping* and *figure-ground segregation*. In grouping, one perceives some elements "as going together more strongly than others." (Wagemans *et al.* 2012: 1180) This allows the system to determine what the *basic units* of perception are. In figure-ground segregation, these units are assigned relative locations and shapes. Thus, if two surfaces share a border, figure-ground segregation determines which surface 'owns' it. The owner of the edge is perceived as the figure and the other as the background (or 'ground' for short). Typically, the figure acquires a determinate shape, while the ground is perceived as shapeless and extending behind the figure (Peterson 2003: 87–90).

Grouping is a necessary condition to individuate objects because objects are, at the very least, groups of elements (Kant 1781/1787). Nevertheless, grouping is not sufficient for reference to *ordinary* objects because entities we do not usually classify as objects are also groups of elements. They include collections like a flock of geese or two-dimensional entities like a set of dots on a screen. Hence, when we say that selection requires individuation, we have in mind a more demanding notion of an object.

A crucial difference between a dot and an ordinary object is that the latter is a three-dimensional entity. Indeed, the figure-ground distinction plays a central role in our ability to parse some portions of reality as three-dimensional entities. To this end, figure-ground discrimination determines the owners of borders. In the visual modality, border assignment is crucial because some borders can correspond to shadows or junctions of planar surfaces differing in reflectance. By contrast, when the borders are assigned to an object, the subject perceives it as relatively closer than the surrounding surface. This relative closeness indicates that depth perception is involved, so the wholes are not merely parsed as two-dimensional entities but as entities located in a three-dimensional layout.

Campbell construes individuation as a relation between a subject (via an object file) and an object in the world (Section 3). It is only when one recognizes the object or makes a further use of a property that representational contents come into the picture. In what follows, I challenge Campbell's relational view of individuation by arguing that

the two kinds of perceptual organization distinguished earlier have primitive representational contents.⁸

4.2 Perceptual Organization and Correctness Conditions

In this sub-section, I suggest that it is natural to associate correctness conditions with episodes of perceptual organization. I hope that everything I will say here will be fairly uncontroversial to most theorists of perception, at least anti-representational singularists like Campbell.

Consider grouping. Intuitively, a visual experience as of a *single* cube differs from a visual experience as of *two* cubes. The difference does not concern the property (Being a cube), which is experienced as instantiated in the two cases. It concerns the *cardinality* of perceptual units that instantiate that property. Similarly, if you are confronted with a cloud of flies, it makes a great difference whether you focus on a *particular* fly or on the *cloud of flies*. Since grouping leads us to parse some elements as more tightly linked to each other than others, it is possible to associate correctness conditions with grouping episodes. Intuitively, a grouping *G* is correct if and only if *G* matches how things are clustered in the world. It is incorrect otherwise.⁹

Consider now figure-ground segregation, illustrated in the Kanizsa and Gerbino (1976) display (Fig. 2). Most observers report the black regions as figures. It is generally accepted that this occurs because convexity can be a reliable figural cue in some contexts (Peterson and Kimchi 2013: 10–1). Interestingly, this point is naturally cashed out in terms of correctness conditions. Suppose that this is a two-dimensional projection of a real-world scene. If the subject sees the black regions as figures, that episode of figure-ground segregation is correct if the black regions are closer to her. It is incorrect otherwise. Besides, if she parses the scene as containing *four* black columns and *one* white background that extends behind them, that episode of figure-ground segregation is correct if there are *four* black columns and *one* white background. It is incorrect otherwise. More generally: A figure-ground distinction *D* is correct if and only if its figures correspond to objects in the world that are closer to the observer and its grounds correspond to more distant surfaces. It is incorrect otherwise.

Finally, it is also possible to imagine perceptual states involving incorrect grouping, figure-ground segregation, and tracking. Consider a scenario matching the case in which someone is approaching you from a long distance (Section 1). In this scenario, however, it is neither a person nor a body that is approaching you. Instead, it is a relatively cohesive swarm of bees that is approaching you at a distance. If the swarm of bees managed to follow a relatively continuous trajectory, it could fool you. In that case, you would take that swarm of bees to be a moving body. This case is naturally characterized as involving incorrect grouping, figure-ground segregation, and tracking.

⁸ If any representational content needs a vehicle to carry it, my arguments can be used to justify the introduction of more primitive representational vehicles than object files.

⁹ For different analyses of correctness conditions, see Pautz (2011), Schellenberg (2014), and Siegel (2010). For my present purposes, these details do not matter.

4.3 The Reference-Fixing Role of Representational Contents

In the previous sub-section, I suggested that it is natural to associate correctness conditions with various forms of perceptual organization. These characterizations are so uncontroversial that most theorists could grant them. In what follows, I want to make a case for the stronger claim that these characterizations are not only descriptively adequate but also explanatorily relevant. To this end, I will submit that correctness conditions are not merely associated with perceptual organization. Instead, the existence of these correctness conditions indicates that perceptual individuation is a primitive form of representing-as.

When one engages in grouping, one represents a multiplicity as *one*. When one draws a figure-ground distinction, one represents the owner of the border as closer and detached from a formless background. If my arguments are plausible, they suggest that these primitive perceptual representational contents are necessary to fix the reference of object files. Thus, *pace* Campbell, representational contents do not merely intervene *downstream* of perception, when the cognitive system has access to the selected items. Representational contents are already necessary *upstream* of object files, to explain their acquisition of a referent.

A good way of determining whether an entity e must be posited to account for the occurrence of an event e' is to figure out whether e' could occur without e . If one fails to see how e' could occur without e , one has a reason to posit the existence of e . To be sure, this strategy does not offer a *conclusive* reason for the existence of e . Yet, this reason constitutes a challenge for those theorists who assume that one can account for e' without positing e .

In the present case, we are interested in the facts that explain the activation of an object file and its acquisition of a referent. A good way of determining whether representational contents are necessary to explain this event is to try to figure out how a singularist could explain the activation of object files without positing representational contents. I want to argue that a highly influential strategy to avoid the introduction of perceptual representational contents does not offer an adequate explanation of the activation of object files. Subsequently, I will submit that positing primitive representational contents underlying perceptual organization allows us to offer a better explanation of reference fixing.

According to the representational approach, perceptual individuation is underwritten by a set of representational contents that enable the system to single out objects. The representational approach uses these contents to analyze perceptual individuation as underwritten by primitive perceptual states that represent objects correctly or incorrectly. In the visual case, these primitive representational contents are necessary to explain the activation of an object file. When these contents are correct, the system successfully individuates an object o and the corresponding object file refers to o . When they are incorrect, an object file becomes active but fails to refer.

The singularist might reject this proposal by providing a disjunctive analysis of perceptual individuation.¹⁰ To this end, she might adapt Martin's (2006) epistemic analysis of hallucination to perceptual individuation. After all, singularists have often developed epistemic analyses of hallucination to provide non-representational accounts

¹⁰ Indeed, Campbell (2002: 130) seems to be attracted by a disjunctivist analysis of perceptual individuation.

of hallucination (Fish 2009). Roughly, these analyses establish a deep asymmetry between the ‘good’ and the ‘bad’ cases, so the latter can be epistemically analyzed in terms of the former. If successful, this approach avoids the introduction of representational contents to analyze perceptual states.

Let us consider an application of the epistemic account to perceptual tracking. A ‘bad case’ of perceptual tracking could be construed by devising a scenario matching the case in which someone is approaching you from a long distance (Section 1). Recall the swarm-of-bees example (Section 4.2). In the imagined case, a relatively cohesive swarm of bees is approaching you at a distance. Hence, you take that swarm of bees to be a moving body. This means that an object file is activated to track the swarm of bees.

Critics of the representational analysis might want to follow Martin’s lead and say: When one keeps track of the swarm of bees in the bad case, one is in a perceptual state in which one cannot know, by reflection alone, whether one is tracking a relatively cohesive collection of entities or a moving body. This analysis would yield a disjunctivist account of the correctness conditions of perceptual tracking.¹¹ Roughly, a tracking episode *E* is correct if there is *one* volumetric object at the other end of the referential relation. A tracking episode *E*’ is incorrect if there is *no* volumetric object or there are *many* volumetric objects at the other end of the referential relation. The epistemic component of the analysis says that the tracking episode *E*’ is indistinguishable by reflection alone from the correct tracking episode *E*.

This analysis would allow us to capture the prior intuitions about the association of correctness conditions with perceptual individuation. Besides, it would allow us to keep Campbell’s reliance on relations to objects as the most fundamental characterizations of reference. Although one could *classify* individuation and tracking episodes as correct or incorrect, this classification would not entitle us to posit representational contents to *explain* perceptual individuation and tracking.

How might an epistemic fact—the indiscriminability between the good and the bad cases—ground the *activation* of an object file in the bad cases? To be sure, if one finds oneself tracking a swarm of bees as a moving body, one’s experience may be indiscriminable from a case in which one is tracking a moving body. Nevertheless, this indiscriminability seems to be a *consequence* of having activated an object file to keep track of the swarm of bees, so it cannot explain it. Let me elaborate.

If a bad case of tracking is indiscriminable from a good case of tracking, both cases must have the same functional role. In the mental file framework, if two tracking episodes have the same functional role, the same type of file realizes them (Section 1). Consequently, if two tracking episodes are indiscriminable, the same kind of file has been activated. In the good case, an object file was activated to keep track of a moving body. In the bad case, an object file was activated to keep track of a swarm of bees. Thus, indiscriminability *presupposes* the activation of an object file. If the two tracking episodes are indiscriminable, it is because the same type of file underwrites them.

This argument hinges on the principle that indiscriminability between the good and the bad cases requires sameness of functional role (Fish 2009). One might reply that this premise is not mandatory, though. After all, functional roles can be broadly individuated. On a broad understanding, the functional roles of the two cases are different, for they involve different worldly entities. The good case has one volumetric

¹¹ A similar analysis could be provided of static forms of perceptual reference.

entity at the other end of the referential relation. The bad case has a multiplicity of entities at the other end of the referential relation.

Although broad functional roles may be relevant in some explanatory contexts, all we need here is the plausible principle that the bad and the good cases share the same narrow functional role ('n-functional role' for short). An n-functional role only considers the relations that a mental state bears to other mental states. Now, if the bad and the good cases were n-functionally different, they would be discriminable. Being n-functionally different, it would be possible—at least in principle—to detect minor differences in how correct or incorrect tracking episodes are related to other mental states. This would make the two states discriminable. Yet, what makes the epistemic analysis so attractive is that, however ideal one's epistemic situation is, the subject cannot discriminate, *even in principle*, the good from the bad cases (Martin 2006).¹²

To sum up: If two tracking episodes are indiscriminable, they share the same n-functional role. In the mental file framework, sameness of n-functional role requires sameness of type of file. As a result, if good and bad tracking episodes are indiscriminable from each other, they share the same type of file. In other words, the activation of the same type of file in the good and the bad cases cannot be explained by their mutual indiscriminability.

It might be objected that accounting for the activation of an object file is only part of the reference-fixing problem. After all, an object file can be activated but fail to refer to a particular object. In the bad case we have been considering, the object file does not refer to the swarm of bees because object files are mental representations analogous to singular terms (Section 1). So it is inappropriate to reject an account of reference fixing just by focusing on the activation of object files in response to a collection of individuals.

Even if one denies that the object file refers to the swarm of bees, the previous example has a more general moral: the disjunctivist is led to offer no account of the activation conditions of object files in the bad cases. This is an unattractive result. Moreover, it is reasonable to assume that the activation conditions of object files are sufficiently similar in the good and the bad cases. It is hard to see how someone who is not antecedently committed to disjunctivism could reject this assumption, given that there is a simpler account that provides a unified explanation of the activation conditions of object files in the good and the bad cases, or so I shall argue.

If perceptual individuation is explained by positing representational contents, the sort of fact that grounds the activation of object files is the same in the two cases. The visual system represents a multiplicity of elements as one volumetric entity. This primitive representation explains why an object file becomes active to track a person approaching from a long distance. Interestingly, the same account can be offered to explain why an object file is activated to track a relatively cohesive swarm of bees. The difference between the two cases lies in the match (or mismatch) of the primitive representations with the world. This is—I take it—the underlying assumption of much contemporary cognitive science. It assumes that the same representational contents underlie the exercise of perceptual abilities in the good and the bad cases. And there

¹² It is common to individuate types of mental files by sameness of n-functional role (Recanati 2012). It is also common to put forward views that guarantee that, if two mental files are of the same type, they have the same n-functional role (Fodor 2008). And this is how cognitive scientists characterize object files (Section 1).

seems to be no good reason to reject this assumption in favor of a disjunctivist analysis of perceptual individuation (see Burge 2005, for a similar line of thought).¹³

5 Representational Singularism and Descriptivism

Singularists often hold that relations to objects are fundamental and unanalyzable. If we are interested in explaining reference fixing, as Campbell is, this radical primitivism is not mandatory. There are more fundamental ways of explaining reference than positing relations to objects. If we want to provide a uniform and informative account of the activation conditions of object files in the good and the bad cases, it is reasonable to introduce representational contents. *Pace* Campbell, representational contents are not only found downstream of perceptual relations. They are also relevant to explain how object files refer to objects.

Some might wonder whether this representational strategy introduces descriptivism back into the picture. I want to suggest, though, that the representational contents of perceptual individuation do not vindicate descriptivism. If my arguments are correct, there is a third way between descriptivism and singularism (as far as visual reference is concerned). I call this view ‘representational singularism.’

Object files represent the instantiation of properties by some objects. They may have the form THIS IS F, where THIS refers to an object and F denotes a property (Pylyshyn 2007; Recanati 2012). Thus, positing more primitive representational contents to explain the activation of object files requires that we introduce different representations underlying the activation of THIS. Since these representations concern pre-objective elements, their contents cannot be carried by predicates of mental demonstratives. The resulting view fits therefore our characterization of the weak acquaintance-first model (WAFM): it posits a non-satisfactional reference-fixing mechanism. Since it departs from the purely causal account of properties in reference fixing, I call it ‘*representational singularism*.’

On the descriptivist model, what fixes reference is the fact that a set of properties is *uniquely* satisfied by the object of reference (Section 2). This means that successful reference must rule out two states of affairs: that the properties have no satisfier and that they have more than one satisfier. Clearly, uses of properties to group items do not fit into this model. Grouping is not a procedure whereby one ascertains whether a property is uniquely satisfied by an object. Instead, it is a means of determining what counts as a perceptual *unit*. To this end, it is not necessary to determine whether there is at least one and at most one object that satisfies a property but rather which elements go together. Hence, grouping is a pre-condition for raising *any* question about the satisfaction of the uniqueness condition. If one asks: ‘Is F uniquely satisfied?’ one must already know what counts as *one* object.

Similarly, in figure-ground segregation, the crucial issue is whether a given border belongs to one or another surface. Border attribution decides which surface acquires

¹³ This argument seeks to show that disjunctivism is not a good strategy for avoiding the introduction of representational contents to explain reference fixing. Yet, it does not undermine *epistemological* disjunctivism. One might still argue that the evidence available in the good cases is not of the same type as the evidence available in the bad cases.

figural status so that it is parsed as a three-dimensional object. Interestingly, the figure is not determined as the unique satisfier of a specific shape. Instead, the figure's possession of a specific shape is the *upshot* of the segregation process. Whereas the figure is perceived as the surface that has a determinate shape, the ground is perceived as a shapeless region that seems to continue behind it. Thus, *the shape is not available independently of the very process of detaching the figure from the background*. So it makes little sense to depict its role as a represented condition in search of a unique satisfier.¹⁴

To be sure, one might re-describe individuation by using the vocabulary of descriptivism. To this end, one might submit that there is at least one relation that must be uniquely satisfied by the pre-objective elements that are represented by the visual system: the relation of belonging to the same object. If this relation is not uniquely instantiated by those elements—so the argument goes—the scene has not been accurately parsed into objects.¹⁵

Although there is a sense in which the unique instantiation of the relation of belonging to the same object is necessary for successful individuation, this is not sufficient to count the present proposal as descriptivist. First, recall that descriptivism holds that the representation of the properties of objects is a *pre-condition* for referring to those objects (Section 2). Yet, there is no reason to believe that the representation of the relation of belonging to the same object is a pre-condition for referring to objects. If the visual system needed a representation of the relation of belonging to the same object in order to individuate objects, the account would be circular. It would presuppose a prior ability to determine what counts as the same object in the scene. Second, although successful reference might require that the relation of belonging to the same object must be uniquely satisfied, this re-description violates the formulation of the satisfactional claim, which follows Russell's theory of definite descriptions. In Russell's model, the reference-fixing properties must be uniquely satisfied by the *object of reference*, and not by the *pre-objective elements* that turn out to be parts of that object (Section 2). This explains why the representational contents necessary for perceptual individuation must be non-descriptivist. Whereas mental file theorists characterize descriptivism as the claim that the properties denoted by the predicates of *object files* fix the reference of those files, the properties used in perceptual individuation cannot be denoted by the predicates of object files. Otherwise, we would be led to conflate the properties of objects with the properties of their parts.

I do not mean to imply that there are no commonalities between descriptivist theories of reference and the representations underlying perceptual organization. The most obvious similarity is that individuation can be characterized by using the expression 'satisfy.' In the Ishihara Color Vision Test, the visual system is sensitive to the

¹⁴ Of course, some shapes must be available before the system distinguishes figure from ground. These shapes are necessary to characterize some organizational principles, such as convexity or symmetry. Nevertheless, these shapes are not attributes of whole objects but of object-parts, as can be seen in Fig. 2. The columns are not entirely convex; only some of their parts are.

The Kanizsa and Gerbino display also shows that some forms of figure-ground distinction exploit a *pattern of similar configurations*. If one reduces the number of black surfaces to one, it is unlikely that one experiences the remaining black surface as a figure. Hence, in some cases, the figure is not determined as the unique satisfier of a given property. Instead, it is determined as *one of the satisfiers* of a given property.

¹⁵ I owe this objection to an anonymous referee of this journal.

satisfaction by some dots of the property ⟨Being green⟩ and of relational properties like ⟨Having a similar hue⟩ and ⟨Being relatively close⟩. Nevertheless, this re-description introduces different semantic relations from those traditionally exploited in the descriptivist literature. The dots are *not the object of reference* but *parts* of that object. This is sufficient to count their relations as non-satisfactional in the relevant sense (Section 2).

It is easy to overlook this point because *some* of the properties we perceptually attribute to objects are also used to individuate them. Think of the ‘5’ in the Ishihara Color Vision Test. In this example, the property ⟨Being green⟩ is both a property of *each* individual dot and a property of the ‘5.’ Nevertheless, this does not hold for all the properties used for individuation. The relational properties ⟨Being relatively close⟩ and ⟨Having a similar hue⟩—which are used for individuation as well—are not properties of the object but of the pre-objective elements that turn out to be parts of that object.

This difference is relevant to the debate between singularism and descriptivism. An influential objection to descriptivism is that it faces circularity problems. If reference requires the determination of the unique satisfier of a property *F*, there must be a prior means of referring to the bearer of that property (Campbell 2009; Pylyshyn 2007). Once we move from objects to pre-objective elements, this objection cannot get off the ground. After all, individuation abilities do not presuppose unexplained abilities to refer to objects. They only presuppose abilities to single out pre-objective elements, which are parsed as individuals when some organizational principles are satisfied.¹⁶

To be sure, the topic of perceptual individuation would require a more detailed analysis. Nevertheless, what I have said here supports the conclusion that perceptual individuation has non-satisfactional contents. These representational contents do not exploit the attributive relations that obtain between objects and properties but rather the relations that obtain between pre-objective elements. Some relational properties are used to draw a line between a whole and its surrounding non-parts.

6 Representational Singularism and Causal Singularism

Most contemporary work on reference assumes that descriptivism and singularism are the only possible accounts of reference fixing. I cast doubt on this assumption by focusing on the visual fixation of reference. The view I put forward emphasizes the existence of non-satisfactional representational contents that underlie perceptual individuation. In this section, I contrast the proposed view with causal singularism.

A common idea associated with singularism is the claim that perceptual relations are fundamental and non-analyzable. Consider a flowchart analysis of Campbell’s account

¹⁶ I do not mean to imply that descriptivism cannot respond to the circularity problems. I only mean that those problems do not even arise for representational singularism. A different circularity problem could be formulated for representational singularism if individuation had to operate over *parts* of objects. If the system had to know beforehand whether what it currently perceives is a part, perception could never get off the ground. In order to be in a position to discover the parts of a scene, one should already have singled out the corresponding object. For this reason, a correct account of perceptual organization must describe the system as having access to portions of the scene before any part/non-part distinction is made. That is precisely why I am using the phrase ‘pre-objective element.’ Parsing some elements as parts is one of the outputs of organizational processes.

(Fig. 3). His model has the interesting feature of making room for a use of properties that differs from their use as the contents of object files. Campbell thinks that a primitive, non-representational mechanism enables the visual system to use a property to select an item in the world. This use for selection is primitive because the system does not have to token any concept of the relevant property, nor perform inductive generalizations involving that property. Unfortunately, his model offers no account of incorrect individuation.

Representational singularism, by contrast, introduces representational contents to explain how the reference of object files is fixed. Contrary to descriptivism, it does not construe those reference-fixing contents as properties denoted by the predicates of object files. And contrary to singularism, it introduces more primitive representations that implement grouping, figure-ground segregation, and possibly other kinds of perceptual organization (Fig. 4).

We lack the empirical knowledge to offer a detailed analysis of the representational contents that underlie perceptual individuation. Yet, the following indications might be instructive. Whereas the causal singularist associates an object with THIS_x via an acquaintance relation, the representational singularist introduces reference-fixing contents defined over sequences of pre-objective elements like dots, lines, surfaces, and so on. These elements are pre-objective because they do not presuppose that individuation has already occurred. They are rather the input that enables visual systems to single out items in the world.

Thus, the reference-fixing content of $|\text{THIS}_x|$ is a function $f(|E_1, E_2, \dots, E_{n-1}, E_n|)$ of a sequence $|E_1, E_2, \dots, E_{n-1}, E_n|$ of pre-objective elements $E_1, E_2, \dots, E_{n-1}, E_n$. In order to form a reference-fixing sequence, the elements should comply with the principles that govern grouping, figure-ground segregation, and other forms of perceptual organization. When the argument of the function is given, THIS_x becomes active. Yet, the activation of THIS_x is not sufficient to fix reference. An activation of THIS_x refers to an object in the world only if there is a sequence $|E_1, E_2, \dots, E_{n-1}, E_n|$ of pre-objective elements $E_1, E_2, \dots, E_{n-1}, E_n$ in the world that matches and stands in an (appropriate) causal relation to the represented sequence.¹⁷

When Campbell (2002, 2009) argues that a relational view is preferable to any representational alternative, his underlying assumption is that the only available representational contents are those expressed by sentences containing singular terms referring to objects, predicates denoting properties, and quantifiers ranging over objects and properties. If my approach is on the right track, however, the relational view is not

¹⁷ This analysis came to my mind after reading Fine's (2007) eye-opening defense of semantic relationism. Contrary to Fine, however, my semantic contents are not defined over sequences of objects but of pre-objective elements.

Our reference-fixing contents also bear a superficial similarity to Kaplan's (1977) characters. Kaplan's characters fix the reference of indexicals but are not part of the content of those indexicals. Similarly, our reference-fixing contents fix the reference of object files but are not part of their contents. Nevertheless, our reference-fixing contents differ from Kaplan's characters. First, Kaplan's characters are functions from contexts to contents. Technically, Kaplan's contexts are n -tuples of parameters. While these n -tuples are typically understood as lists of relatively independent parameters, the arguments of our reference-fixing contents are sequences of parameters that are related to each other by organizational principles. Second, whereas Kaplan's contexts are often construed as sets of *individuals* or *particular* times or locations, the arguments of our reference-fixing contents are more primitive than individuals. They are pre-objective elements that serve as input to determine the individuals present in a scene.

compulsory. If perceptual reference is fixed by *pre-objective* representational contents, relations to objects have been elucidated by looking at the representational structures of grouping, figure-ground segregation, and other kinds of perceptual organization. In the visual modality, it is because we can represent some elements in the scene as parts of one volumetric entity that we can stand in epistemically rewarding relations to objects.¹⁸

Some readers might protest that representational singularism does not really stop the regress faced by descriptivism. After all, the ability to single out pre-objective elements leaves a number of questions open. If the pre-objective elements are represented, how is the reference of these representations fixed? Is it fixed via a purely causal mechanism?¹⁹

I admit that representational singularism does not offer a final solution to the regress problem. Nevertheless, it does constitute progress in relation to descriptivism and causal singularism. To show why, let me distinguish two kinds of homuncular theories. *Circular* theories try to explain an ability *A* by postulating another ability *B* that requires the exercise of *A*. Traditional forms of descriptivism are circular in this sense because they explain reference by postulating abilities that require the exercise of the very ability to refer to objects. *Non-circular* theories, by contrast, try to explain a complex ability as the interplay of simpler abilities that are taken as primitive. As Dennett (1978: 123–4) famously argued, psychological explanations are homuncular in this second sense:

If one can get a team or committee of *relatively* ignorant, narrow-minded, blind homunculi to produce the intelligent behavior of the whole, this is progress. A flowchart is typically the organizational chart of a committee of homunculi (investigators, librarians, accountants, executives); each box specifies a homunculus by specifying a function *without saying how it is to be accomplished* (one says, in effect: put a little man in there to do the job). (Dennett 1978: 123–4).

Representational singularism is homuncular in Dennett's sense. It factorizes the ability to refer to objects as the interplay of the abilities (1) to single out pre-objective elements and (2) to organize those elements by principles of grouping, figure-ground segregation, and so on.

How is the reference to these pre-objective elements fixed? This is an empirical issue that cannot be settled from the armchair. My view is very close to Dennett's: the regress will come to an end when we get a flowchart analysis that lands us with homunculi so stupid that they could be replaced by a machine. Thus, the regress will come to an end when we identify a purely causal process. A number of theoretical options are open. One might claim that the finest-grained flowchart analysis would land us with *atomic* representations that refer to pre-objective elements in a purely causal manner. I am not sure whether the devices that single out pre-objective objects are atomic representations in that sense. One might also claim that the finest-grained flowchart analysis would

¹⁸ A multisensory model of reference fixing ought to make room for other reference-fixing mechanisms, such as the auditory perception of events. On the plausible assumption that voices have the function of enabling humans to individuate and recognize their conspecifics (Matthen 2010), listening to a voice may suffice to open an object file or even a recognitional file (if one is familiar with the voice).

¹⁹ These questions were prompted by a referee of this journal.

land us with basic parts of the functional architecture that cannot be plausibly described as representations. These are difficult topics that ultimately depend on one's choice of a functional architecture. They are at the heart of the controversy between symbolic, connectionist, and dynamic architectures. Although these questions are largely open, my main point remains: representational singularism fares better than descriptivism and causal singularism because it factorizes perceptual reference as the interplay of simpler abilities.

7 Representational Singularism and Existentialism

The focus of this paper was Campbell's picture of the causal use of properties to individuate objects. One might wonder, however, whether representational singularism is strictly necessary to disprove Campbell's anti-representational account. After all, there are other theories that also introduce contents to characterize visual experience while avoiding the shortcomings of descriptivism. A good example is the existential account of perceptual content (McGinn 1982; Davies 1992; Jackson 2012). On this view, visual experiences represent objects in general terms, i.e., as the satisfiers of some properties. If a subject has an experience as of a red cube in front of her, the content of the experience could be conveyed by the sentence: 'There is a red cube at *L*.' This reliance on an object's satisfaction of properties makes the existential view very similar to descriptivism. Nevertheless, the existential view drops the uniqueness condition. There are two reasons for this move: "an object may fit the general conditions comprised in the perceptual mode of presentation and yet not be perceived, and the object perceived may fail to fit the content of the experience." (McGinn 1982: 53) Consequently, defenders of the existential view claim that the objects of perception are not fixed by the general contents of perception but by (appropriate) causal relations.

How does representational singularism compare to existentialism?²⁰ I shall argue that the alleged advantages of existentialism over descriptivism can be easily accommodated by representational singularism. This will lend additional support to the claim that representational singularism is not a version of descriptivism (Section 5). Moreover, the existential views available in the literature offer a less satisfactory explanation of reference fixing than the one provided by representational singularism.²¹

An alleged advantage of existential views is that they accommodate an important phenomenological fact: numerically different objects can be experienced as numerically the same (Davies 1992: 26; McGinn 1982: 60). Let us term it the 'indiscriminability intuition.' In order to capture it, defenders of existentialism ban singular terms from any characterization of the content of visual experiences.

Although representational singularism is not motivated by phenomenological considerations, nothing prevents it from accommodating the indiscriminability intuition. To this end, one might *functionalize* the sequences of pre-objective elements that constitute the arguments of reference-fixing contents. Let me illustrate this move by means of the

²⁰ This question was prompted by an anonymous referee of this journal.

²¹ So-called veridical hallucinations and illusions have also played a central role in debates on existential accounts of perceptual content. These cases raise difficult questions I cannot adequately examine in this paper. I will simply follow Tye (2011) and claim that, *if* veridical illusions and hallucinations constitute a legitimate test for theories of perceptual content, the problem they pose is also a problem for existential views.

‘5’ in the Ishihara Color Vision Test. One might identify the argument of the reference-fixing function with a *particular* sequence of dots. This view predicts that numerically different sequences lead to the activation of numerically different object files. Nevertheless, this interpretation is not mandatory. Indeed, one could claim that a reference-fixing function can lead to the activation of the same object file in response to *any* sequence of entities that satisfies the same individuation principles and properties as the particular dots in Fig. 1. In other words, the reference-fixing function can activate the same object file to represent a ‘5’ when qualitatively identical sequences of pre-objective elements are presented to the subject. As a result, if two numerically different arrangements of dots are experienced as one and the same ‘5,’ the reference fixing mechanisms that led to the activation of the same object file in the two cases possibly operated over qualitatively identical sequences of pre-objective elements.²²

Another alleged advantage of existentialism is that it overcomes an important problem of descriptivism. As McGinn (1982: 53) points out, “the object perceived may fail to fit the content of the experience.” This is well illustrated by the case in which one misperceives a stick half-immersed in water as bent. In this scenario, although the stick is not the unique satisfier of the property ⟨Being bent⟩, one’s perceptual state is still about the stick. Since the existential view drops the uniqueness condition, it can avoid this problem. Indeed, it accommodates this case by supplementing the general content of perception with an appropriate causal relation.

Representational singularism can accommodate this fact because it endorses an amended version of Campbell’s two-fold distinction between (1) using a property to individuate an object and (2) verifying a proposition to the effect that the object has a property. Once this distinction is introduced, the problem faced by perceptual descriptivism vanishes. If I can misperceive a straight stick as a bent stick, it is because the property ⟨Being bent⟩ was not used by my visual system to individuate the stick. This property was just *misattributed* by the visual system to the stick. In the mental file framework, this means that the property ⟨Being bent⟩ is represented by a predicate in the object file. Before this (or any other property) can be ‘written’ in the object file, there must be organization mechanisms that fix reference. According to representational singularism, these organization mechanisms involve primitive representational contents that exploit other of properties instantiated by sequences of pre-objective elements.²³

To sum up, the two main advantages of existentialism over descriptivism can be easily accommodated within representational singularism: the latter has the resources to accommodate the indiscriminability intuition and to explain successful perceptual reference in cases of perceptual misattribution. This confirms the previous claim that representational singularism is not a version of descriptivism (Section 5). My final

²² The reason I wrote ‘possibly’ is that the account should be qualified with appropriate caveats concerning psychological limitations and contextual factors. Perceptual and memory limitations, distraction or even different task demands might lead a subject to activate numerically different object files in response to qualitatively identical sequences of pre-objective elements and the same object file in response to qualitatively different sequences of pre-objective elements.

²³ If causal relations only obtain between objects as possessors of some properties, holders of existential views should avail themselves of a similar distinction to solve the problem. They should carefully distinguish the properties that enter into the reference-fixing causal relation from the properties that enter into the general content of perception.

point is that the present account of reference fixing is superior to the account offered by existential views.

Recall the indiscriminability intuition. Imagine a subject who has two experiences of two numerically different but indiscriminable objects at two different times. On the plausible assumption that the two experiences have different referents, the general content of those visual experiences is not sufficient to fix their reference. That is another reason why most defenders of existential views supplement the general content of visual experiences with causal relations to objects (Davies 1992: 26; Jackson 2012: 204–5; McGinn 1982: 53).

I agree that causation is necessary to fix reference (Section 6). Nevertheless, I have some qualms with the use of causation in the existential framework. First, a highly influential version of this view construes the causal relation as represented in the content of perceptual states (Searle 1983: 48, 123). On this approach, a visual experience as of a red cube has a two-tiered content: (1) that there is a red cube and (2) that that red cube stands in such and such causal relation to this very experience. Unfortunately, this approach is psychologically and phenomenologically implausible (Tye 2011: 176). Second, even if one managed to respond to this objection, the resulting picture of reference fixing would not be very informative. Indeed, it is more informative to claim that some organization processes underwritten by mental representations fix visual reference—as representational singularism does—than just to say that it is fixed by an appropriate causal relation. Third, causal models of reference fixing lack—while representational singularism has—an account of the activation conditions of object files in cases of incorrect individuation (Section 4.3). For these reasons, representational singularism constitutes a decisive step forward in the project of explaining reference fixing.²⁴

²⁴ In his recent work on perception, Tyler Burge sketched an analysis of reference that does not fit into the singularist/descriptivist dichotomy. According to Burge (2010: 464), “to represent something as a body, the individual’s perceptual system must segment a three-dimensional whole from a surround by [...] principles for identifying cohesiveness and boundedness of three-dimensional volume shapes.” What is the difference between Burge’s account and representational singularism?

We offer different analyses of the contents that fix reference to objects. Whereas representational singularism posits a reference-fixing function defined over a sequence of pre-objective elements, Burge holds that reference to objects is fixed by noun phrase structures of the form THIS F. Burge’s analysis is partly supported by his Second Thesis on *De Re* States: each demonstrative representation “must be associated with a nonschematic attributive [...] that guides the singular representation.” (Burge 2009: 275) If representational singularism is correct, Burge’s Second Thesis is mistaken. Indeed, representational singularism shows that the reference of an object file need not be ‘guided’ by any attributive associated with the object file. After all, its reference can be fixed by *other* representational structures that lead to the activation of the object file. Thus, an object file can become active before (and independently of) any attribution of properties to the referred object.

A good way of illustrating the difference between the two views is by reflecting on the example of the stick half-immersed in water. According to representational singularism, reference is not undermined by the misattribution of the property (Being bent) because the properties that fix reference are not the same as the properties the visual system attributes to objects. In other words, there is a principled difference between individuating objects and attributing properties to those objects. Given his commitment to the Second Thesis, Burge is compelled to treat object individuation on the model of property attribution. Thus, on Burge’s view, when there is reference to objects, there must be some attributives that are veridical of *those objects* (Burge 2009: 293–5; 2010: 151). Unfortunately, Burge offers no empirical support for this bold claim. I think this is a mistake inherited from descriptivism. As was pointed out in Section 5, many properties used for individuation are not properties of the referred object but of the pre-objective elements that turn out to be parts of that object.

Acknowledgments I presented earlier versions of this paper at the Instituto de Investigaciones Filosóficas-UNAM (7th May 2014), the *5th Workshop on Language, Cognition, and Context* (23–25 May 2014), the University of Antioquia (29th May 2014), and the University of Geneva (25th September 2014). I am grateful to the audiences for their comments, especially to Santiago Amaya, Julien Deonna, Richard Dub, and Fabrice Teroni. I am also indebted to Santiago Arango, Reinaldo Bernal, Ariel Cecchi, Jérôme Dokic, Gregory Bohner, and an anonymous referee for their comments on earlier drafts of this paper. This work was funded by the Swiss National Science Foundation (research grant No. 100012-150265/1).

Compliance with Ethical Standards The present work has not been published before and is not under consideration for publication anywhere else. I declare that there is no actual or potential conflict of interest including any financial, personal or other relationships that could inappropriately influence, or be perceived to influence, my work.

Appendix

Fig. 1




Fig. 1 If you are a normal trichromat and are not looking at this pattern in *black* and *white*, you will see a green numeral 5

Fig. 2



Fig. 2 Figure-ground segregation by convexity. The *black* and *white* regions have the same area. Yet, the visual system parses the *black* regions as figures and the *white* ones as ground. Reproduced from Peterson and Kimchi (2013)

Fig. 3




Fig. 3 Campbell's selection/access model of reference fixing. **a** A causal use of a property enables the system to select an object or region. The quotation mark indicates that Campbell's model offers no account of incorrect individuation. **b** Once the object or region is selected, the system can represent its properties. These representations may be answers to queries like 'WHAT COLOR IS THIS?'

Fig. 4




Fig. 4 Representational singularism. **a** Principles of perceptual organization are implemented by primitive representations. **b** The interaction among those representations yields a representation of various elements as one volumetric object. This representation explains the activation of object files in the good and the bad cases. Only when the representation matches the world and stands in an appropriate causal relation to the represented sequence, the object file refers to a determinate object

References

- Bach, K. 1987. *Thought and reference*. Oxford: Clarendon.
- Block, N. 1995. On a confusion about a function of consciousness. *Brain and Behavioral Sciences* 18(2): 227–247.
- Brewer, B. 2011. *Perception and its objects*. Oxford: Oxford University Press.
- Burge, T. 1977. Belief *De Re*. *The Journal of Philosophy* 74(6): 338–362.
- Burge, T. 2005. Disjunctivism and perceptual psychology. *Philosophical Topics* 33(1): 1–78.

- Burge, T. 2009. Five theses on *De Re* states and attitudes. In *The philosophy of David Kaplan*, ed. J. Almog and P. Leonardi, 246–324. Oxford: Oxford University Press.
- Burge, T. 2010. *Origins of objectivity*. Oxford: Clarendon.
- Campbell, J. 2002. *Reference and consciousness*. Oxford: Clarendon.
- Campbell, J. 2007. The metaphysics of perception, *Philosophical Issues 17: The Metaphysics of Epistemology*: 1–15.
- Campbell, J. 2009. Consciousness and reference. In *The Oxford handbook of philosophy of mind*, ed. B.P. McLaughlin, A. Beckermann, and S. Walter, 648–662. Oxford: Clarendon.
- Campbell, J. 2011. Visual attention and the epistemic role of consciousness. In *Attention: Philosophical and psychological essays*, ed. C. Mole, D. Smithies, and W. Wu, 323–341. New York: Oxford University Press.
- Campbell, J. 2012. Perceiving the intended model. In *Perception, realism, and the problem of reference*, ed. A. Raftopoulos and P. Machamer, 96–122. Cambridge: Cambridge University Press.
- Campbell, J. 2013. Susanna Siegel's the contents of visual experience. *Philosophical Studies* 163: 819–826.
- Carey, S. 2009. *The origin of concepts*. New York: Oxford University Press.
- Chastain, C. 1975. Reference and context. In *Language, mind, and knowledge*, ed. K. Gunderson, 194–269. University of Minnesota Press: Minneapolis.
- Davies, M. 1992. Perceptual content and local supervenience. *Proceedings of the Aristotelian Society* 92: 21–45.
- Dennett, D.C. 1978. *Brainstorms: Philosophical essays on mind and psychology*. Cambridge: MIT Press.
- Dickie, I. 2010. We are acquainted with ordinary things. In *New essays on singular thought*, ed. R. Jeshion, 213–245. Oxford: Clarendon.
- Fine, K. 2007. *Semantic relationism*. Malden: Blackwell.
- Fish, W. 2009. *Perception, hallucination, and illusion*. New York: Oxford University Press.
- Fodor, J.A. 2008. *LOT 2: the language of thought revisited*. New York: Oxford University Press.
- Huang, L., and H. Pashler. 2007. A Boolean map theory of visual attention. *Psychological Review* 114: 599–631.
- Jackson, F. 2012. Michael Tye on perceptual content. *Philosophy and Phenomenological Research* 84(1): 199–205.
- Kahneman, D., A. Treisman, and B. Gibbs. 1992. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 24: 175–219.
- Kanizsa, G., and W. Gerbino. 1976. Convexity and symmetry in figure-ground organization. In *Art and artefacts*, ed. M. Henle, 25–32. New York: Springer.
- Kant, I. (1781/1787) *Kritik der reinen Vernunft*. Hamburg: Felix Meiner. 1998.
- Kaplan, D. 1977. Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In *Themes from Kaplan*, ed. J. Almog, J. Perry, and H. Wettstein, 481–563. New York: Oxford University Press.
- Kennedy, M. 2007. Visual awareness of properties. *Philosophy and Phenomenological Research* 75(2): 298–325.
- Kripke, S.A. 1980. *Naming and necessity*. Cambridge: Harvard University Press.
- Martin, M.G.F. 2006. On being alienated. In *Perceptual experience*, ed. T. Gendler and J. Hawthorne, 354–410. Oxford: Oxford University Press.
- Matthen, M. 2010. On the diversity of auditory objects. *Review of Philosophy and Psychology* 1: 63–89.
- McGinn, C. 1982. *The character of mind*. Oxford: Oxford University Press.
- Pautz, A. 2011. What are the contents of experiences? In *The admissible contents of experience*, ed. K. Hawley and F. Macpherson, 114–138. Malden: Wiley-Blackwell.
- Peterson, M.A. 2003. On figures, grounds, and varieties of surface completion. In *Perceptual organization in vision: Behavioral and neural perspectives*, ed. R. Kimchi, M. Behrmann, and C. Olson, 87–116. Mahwah: LEA.
- Peterson, M.A., and R. Kimchi. 2013. Perceptual organization in vision. In *Handbook of cognitive psychology*, ed. D. Reisberg, 9–31. New York: Oxford University Press.
- Pyllyshyn, Z.W. 2007. *Things and places: How the mind connects with the world*. Cambridge: MIT Press.
- Recanati, F. 2012. *Mental files*. New York: Oxford University Press.
- Russell, B. 1905. On denoting. *Mind* 14: 479–493.
- Russell, B. 1912. *The problems of philosophy*. New York: Oxford University Press.
- Schellenberg, S. 2014. The relational and representational character of perceptual experience. In *Does perception have content?* ed. B. Brogaard. New York: Oxford University Press.
- Searle, J.R. 1983. *Intentionality*. Cambridge: Cambridge University Press.

- Shoemaker, S. 1980. Causality and properties. In *Time and cause*, ed. P. Van Inwagen, 109–135. Dordrecht: Reidel.
- Siegel, S. 2010. *The contents of visual experience*. Oxford: Oxford University Press.
- Spelke, E.S. 1990. Principles of object perception. *Cognitive Science* 14: 29–56.
- Tye, M. 2011. The admissible contents of visual experience. In *The admissible contents of experience*, ed. K. Hawley and F. Macpherson, 172–193. Malden: Wiley-Blackwell.
- Wagemans, J., J.H. Elder, M. Kubovy, S.E. Palmer, M.A. Peterson, M. Sinigh, and R. von der Heydt. 2012. A century of gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin* 138(6): 1172–1217.