



Mindless accuracy: on the ubiquity of content in nature

Alex Morgan¹ 

Received: 28 July 2018 / Accepted: 2 November 2018
© Springer Nature B.V. 2018

Abstract

It is widely held in contemporary philosophy of mind that states with underived representational content are *ipso facto* psychological states. This view—the *Content View*—underlies a number of interesting philosophical projects, such as the attempt to pick out a psychological level of explanation, to demarcate genuinely psychological from non-psychological states, and to limn the class of states with phenomenal character. The most detailed and influential theories of underived representation in philosophy are the *tracking theories* developed by Fodor, Dretske, Millikan and others. Tracking theorists initially hoped to ‘naturalize’ underived representation by showing that although it is distinctively psychological it is not *irreducibly* so, yet they ended up developing theories of representation that by their own lights don’t pick out a distinctively psychological phenomenon at all. Burge (*Origins of objectivity*, Oxford University Press, Oxford, 2010) sets out to develop a theory of underived representation that does pick out a distinctively psychological phenomenon. His theory promises to vindicate the Content View and the various philosophical projects that depend on it. In this paper I argue that Burge’s theory *dementializes* representation for the same reason tracking theories do: These theories hold that representations are states with underived accuracy conditions, yet such states are found in all sorts of mindless systems, like plants.

Keywords Accuracy conditions · Circadian clocks · Intentionality · Mental representation · Objectivity · Perceptual constancy

1 Introduction

Following Brentano (1874 [1995]), many philosophers have thought that what distinguishes mental from non-mental phenomena is *intentionality*. Intentionality is now standardly characterized in terms of representational content, or the possession of

✉ Alex Morgan
alex.morgan@rice.edu

¹ Department of Philosophy, Rice University, 6100 Main MS-14, Houston, TX 77005, USA

accuracy conditions. All sorts of non-mental phenomena have representational content, like sentences or pictures. But if we help ourselves to a distinction between entities whose content is *derived* from the interpretative practices of agents, and entities with *underived* content, we can reformulate Brentano's thesis in the contemporary idiom by saying that states with underived representational content are *ipso facto* mental or psychological states. Call this the *Content View*. The Content View is widely assumed in philosophy—so widely assumed, I think, that its influence often goes unrecognized. But there are several places in the literature where commitment to the Content View is more or less explicit.

For example, the Content View plays an important role in debates over Representationalism, the thesis that psychological processes are explained by appealing to the manipulation of internal representations. In a classic articulation of Representationalism, Fodor and Pylyshyn (1988) argue that what distinguishes minded subjects from mindless, reactive systems is that only minded subjects exhibit behavior that is explained at the 'cognitive level' of explanation, a level of analysis at which "states of the system... encode properties of the world" (p. 10). Similarly, Adams and Aizawa (2008) seek to defend a traditional form of Representationalism from the heterodox 'extended mind' hypothesis, which holds that a subject's psychological states might be partly realized by derived representations in her environment, like notebooks that serve as a kind of external memory (Clark and Chalmers 1998). Adams and Aizawa insist that psychological states and processes necessarily involve the manipulation of states with *underived* representational content, and that such states are in fact only found inside a subject's mind or brain. All of these philosophers endorse mainstream views on which contents are understood in terms of accuracy conditions, so they seem committed to the view that if a state has underived accuracy conditions it qualifies as psychological.

One also finds commitment to the Content View in debates over Intentionalism, the thesis that the phenomenal character of a conscious state is fixed by its representational content. Clearly not all entities with representational content qualify as psychological states, let alone states with phenomenal character—witness words or pictures. So a central question for the Intentionalist concerns what it takes for a representational state to qualify as phenomenal (Kriegel 2002). Tye (2000) develops a form of intentionalism on which a representational state has phenomenal character when it is 'poised' to interact with a subject's beliefs, desires and other psychological states. Tye doesn't say what it takes for a state to qualify as psychological, but if an Intentionalist of his stripe were to endorse one of the standard forms of Representationalism just bruted, she too would presumably be committed to the Content View. Other forms of Intentionalism lead more directly to the Content View. For example, Bourget (2010) seeks to avoid the philosophical burden of articulating precisely what it takes for a representational state to be 'poised' by holding that representational states have phenomenal character, and hence presumably qualify as mental states, simply by virtue of having *underived* content—where, again, content seems to be understood standardly, in terms of accuracy conditions.

So there are several interesting and important philosophical projects that seem to presuppose the Content View, the view that states with underived accuracy conditions are psychological in nature. The task of this paper is to argue that the Content View is

false. I'll argue that many *mindless* systems have internal states with underived content or accuracy conditions. Since there is no canonical theory of underived content in philosophy, my argument will be somewhat indirect and inductive. I'll examine the most mature and influential theories of underived content in the literature and draw general lessons. I begin in Sect. 2 by looking at a family of theories that can lay the best claim to canonical status, the 'tracking theories' found in the naturalizing semantics tradition, which characterize content in terms of the environmental conditions that would have to be satisfied for an information-bearing state to fulfill its biological function. As I'll explain, tracking theories were initially intended to naturalize the representational content of mental states, but in doing so they end up ascribing content to all sorts of mindless systems, like plants. That is, tracking theories *dementalyze* representation.

However, the main focus of this paper is a theory of representation developed by Tyler Burge (2010), which is explicitly intended to avoid the dementalizing consequences of tracking theories. As I explain in Sect. 3, Burge argues that if we attend to successful explanations of perceptual constancy capacities in perceptual psychology, we see that ascribing accuracy conditions to perceptual states is explanatorily indispensable. On these grounds Burge holds that perceptual states comprise the most fundamental kind of psychologically distinctive representation, a kind of representation that is likely irreducible to the kind of non-psychological information registration picked out by tracking theories. Burge's theory is the most detailed, richly argued, and influential theory of distinctively psychological representation currently on offer in contemporary philosophy, so it offers the best hope for someone who wishes to defend the Content View and the various philosophical projects that depend on it.

In Sects. 4, 5 I argue that Burge's theory in fact dementalizes content for the same reasons tracking theories do. My objection takes the form of a dilemma. Burge holds that perceptual states comprise a psychological kind of representation because ascribing accuracy conditions to perceptual states plays an indispensable role in successful psychological explanation, but he doesn't provide a principled account of *why* states with explanatorily indispensable accuracy conditions are psychological in nature. On a natural interpretation of his view, perceptual states qualify as psychological because they are formed through the exercise of perceptual constancy capacities that are themselves inherently psychological. I foreclose this interpretation by showing that constancies as Burge understands them are found in uncontroversially mindless systems, like digital cameras. Burge's considered view seems to be that states with explanatory accuracy conditions are psychological because accuracy conditions in fact only play an explanatory role in psychology. I argue to the contrary that biologists appeal to the accuracy conditions of circadian clocks to explain the adaptive behavior of plants. Circadian clocks in plants thus seem to satisfy the conditions of Burge's view of representation, yet plants presumably don't have minds. So Burge and tracking theorists dementalize representation for the same underlying reason: There is nothing distinctively psychological about explanatory accuracy.

This is not to deny that there is a real and important distinction between the intentional states of psychological subjects and the informational tracking states of plants; it is just to deny that underived accuracy conditions capture this distinction. In Sect. 6, I close the paper by diagnosing why the focus on underived accuracy inevitably

dementializes representation, and offer a prescription for *rementalizing* representation. I suggest that philosophers have been apt to regard states with underived accuracy conditions as psychological because they have been insufficiently attentive to distinct senses of ‘internal’ and ‘objective’. States with underived accuracy conditions are ‘internal’ in the sense of being internal *parts* of a system, but not necessarily in the sense of contributing to a subject’s *inner perspective* on the world. They are ‘objective’ in the sense that they track entities that are *in fact* mind-independent, but not in the sense of presenting entities *as* mind-independent to a subject. I close by suggesting that if we are to capture a distinctively psychological kind of representation, we might try looking at the neural mechanisms that integrate multimodal information into a unified spatial framework such that the information is accessible and attributable to a whole subject, and presents objects as standing in ‘objective’ spatial relations to the subject.

I should say a few words about methodology before moving on. In asking whether underived representation is inherently psychological, neither I nor my interlocutors must assume that psychological phenomena are demarcated from mindless phenomena by a strict set of necessary and sufficient conditions. I for one think that the distinction between mentality and mindlessness will turn out to be vague, allowing for a range of intermediate cases between paradigmatically minded and mindless systems. A comparison with the concept of life might be instructive. We now have a rich and detailed understanding of how living organisms differ from inorganic matter, but this needn’t be structured around a precise definition that delivers decisive verdicts about whether, say, viruses are alive (Bechtel 2007a; Diéguez 2013). Now, it might in principle have turned out that there is *no* distinction in nature, not even a fuzzy one, that corresponds to our pre-theoretical notion of life. Some eliminativists think that this will be true of the concept of mind. For example, Rorty (1979) claims that ‘mind’ applies to a congeries of unrelated phenomena, and serves primarily as an evaluative, honorific term. I agree that evaluative considerations can exert a tacit and pernicious influence on our thinking about the relation between representation and mind, which is why I employ the clumsy but purely descriptive term ‘dementialize’, rather than Burge’s (2010) ‘deflate’ or ‘debase’. But whether there is a distinction in nature between minded and mindless systems cannot be decided from the armchair. Like others interested in the relation between representation and mind, I assume as a tentative working hypothesis that there is such a distinction, which we can illuminate by honing theories against paradigmatic cases. Let’s see how far that assumption can take us.¹

2 Tracking theories

Analytic philosophers have long been suspicious of the naturalistic *bona fides* of intentionality. Rightly or wrongly, many have interpreted Brentano (1874 [1995]) as holding that intentionality is irreducibly mental, and that the essential feature of inten-

¹ Thanks to an anonymous referee for emphasizing the importance of addressing these methodological points. See also note 6.

tional states is that they can represent things that don't exist. It was widely assumed that the most basic form of 'intentional inexistence' involves representational failure, inaccuracy, or *misrepresentation*.² Thus the project of naturalizing intentionality took the form of specifying conditions under which a representation is accurate or correct, using unproblematic concepts from the natural sciences, without presupposing the existence of an agent who interprets the representation. It was hoped that such an analysis would show that the underived content of mental states is part of the natural world after all.

The most detailed and influential theories that emerged from this project differ in detail but share an overall family resemblance. Most fundamentally, they characterize representations as 'internal' states that stand in a causal or informational relation to a specific object, property or state of affairs in the external world. Of course not just any causal relation will do, for a representation might be caused by all sorts of things that it doesn't represent. A well-placed electrical zap to my inferotemporal cortex might elicit an experience as of a face (Afraz et al. 2006), yet if my experience represents anything, it surely represent a face, not the zap. To explain this it was thought that we must appeal not to what a representation is *actually* caused by, but what it is in some sense *supposed* to be caused by. The central agenda of the project of naturalizing intentionality was thus to explain this normative constraint in naturalistic terms, and the prevailing strategy was to appeal to a teleological notion of biological function. The most influential theories emerging from this tradition held, roughly, that the content of a representation is determined by the entity in the external world that it has the function of responding to, or *tracking* (Dretske 1988; Fodor 1990; Millikan 1984). These theories have been aptly dubbed 'tracking' theories.³

Tracking theories notoriously encompass phenomena that don't seem representational in any interesting sense. For example, some philosophers have worried that Millikan's teleosemantic theory entails that "interactions between trees can have content attributed to them" (Allen and Hauser 1993, p. 88), or that "saliva represents food" (Sterelny 1995, p. 256). Others have worried that Fodor's asymmetric-dependence theory might end up ascribing content to pigeon droppings (Adams and Aizawa 1994). Yet presumably neither tree-interactions, saliva, nor pigeon droppings qualify

² Note that the phenomenon of representing something that doesn't exist is often conflated with misrepresentation. This is a mistake. A state might be *about* something that doesn't exist without thereby misrepresenting it. One might desire to meet Elvis without misrepresenting the rhinestoned crooner. Desires exhibit aboutness but just aren't in the business of representing how the world is. Even beliefs, which are in the business of representing the world, might be directed at non-existents without representational failure; consider the true belief that Elvis is dead.

³ There are of course important differences between the theories of Dretske (1988), Fodor (1990), and Millikan (1984). I deliberately abstract away from these differences to reveal what I take to be important underlying commonalities. I lack the space to fully justify my grouping of these theories together, but I note that this classification is now common in the literature (e.g. Burge 2010; Kriegel 2012; Mendelovici 2013). The two most serious objections to this classification are as follows: First, one might point out that Fodor famously argues against appealing to biological function in the context of naturalizing intentionality. This is true as far as his theory of content goes, but Fodor (1990) in fact allows that biofunction might play a central role in explaining other aspects of representation, such as the capacity for misrepresentation. Second, one might point out that Millikan (1984) denies that information or indication plays a direct role in fixing content. Nevertheless, Shea (2007) argues convincingly that Millikan must appeal to an informational 'input condition' to avoid well-known objections to her view.

as representations. Tracking theorists themselves invite the charge that their theories overgeneralize. One of the central examples discussed in the naturalizing semantics literature is the *magnetosome*, an organelle found in certain marine bacteria containing tiny magnetic crystals that help their hosts travel along geomagnetic fields to reach their preferred conditions of deoxygenated water (Blakemore and Frankel 1981). Tracking theorists entered into fierce debate about precisely what magnetosomes represent—Geomagnetic north (Jacob 1997)? The direction of deoxygenated water (Millikan 1989)?—but generally took it for granted that magnetosomes represent something. Yet it's not clear why we should think that. Magnetosomes don't *do* anything obviously representational; all they do is align bacteria along magnetic fields.

To diagnose why tracking theories run the risk of overgeneralizing, we should note that such theories are 'officially' theories of what fixes the content of a representation, not theories of what makes something a representation in the first place. Yet many tracking theorists seem to tacitly assume that their theories of representational content will do double-duty as theories of representational *vehicles*. Thus tracking theories are taken to 'fix the content' of all sorts of things that aren't representational in nature, things to which the ascription of content is gratuitous (Adams and Aizawa 1994; Ramsey 2007; Sterelny 1995).

A full account of what makes something a representation is beyond the scope of this paper. It is enough for present purposes to point out that a representation is something to which the ascription of content is *not* gratuitous. There is wide philosophical agreement that representations constitutively play a distinctive kind of functional role; they are used as surrogates or stand-ins for what they represent, such that *how* they are used can only be fully understood in terms of *what* they represent (Bermúdez 1995; Ramsey 2007; van Gelder 1995). That is, ascribing content to a genuine representation is not gratuitous, but is crucial for explaining what the representation does. To understand how magnetosomes contribute to magnetotaxis, we needn't ascribe content or accuracy conditions to them; it is enough to simply invoke their magnetic properties. Things might be different if magnetosomes could occupy different orientations within bacteria such that their specific orientation were 'read' by control mechanisms that would then drive their bacterial hosts to occupy the same orientation as the magnetosome. We might then have reason to conceptualize the magnetosome as a simple *model* of the geomagnetic field, and explain its contribution to magnetotaxis in terms of its being accurate with respect to the geomagnetic field. But that's not how magnetosomes actually work; they contribute to magnetotaxis by *pulling*, not *modeling*.

Though tracking theorists have generally paid little attention to the explanatory constraints on content ascription, Dretske (1988) is an important exception. His theory of content is rooted in the idea that the specific content of a representation must be relevant to what the representation does: "The fact that [representations] have a content, the fact that they have a semantic character, must be relevant to the kind of effects they produce" (p. 80). Developing this idea, Dretske holds that a state or sub-system *A* of system *S* represents entity *B* when *A* indicates the presence of *B*, and *A* is recruited, either by natural selection or individual learning, to help control *S*'s interactions with *B*, *in virtue of the fact* that *A* indicates *B*. When these conditions are satisfied, the fact that *A* indicates *B* is crucial for understanding the distinctive role that *A* plays within *S*. *A* is used within *S* as a kind of proxy for *B*, and thus plays an intelligibly

representational role. Magnetosomes do not play such a role, hence they are excluded from the scope of Dretske's theory.

Other tracking theorists can and arguably do appeal to similar 'use-conditions' to constrain the scope of their theories.⁴ But here we reach the crucial point. Even if tracking theories were constrained with a use-condition such that they would exclude *non-representational* phenomena like magnetosomes, salivation, and pigeon droppings, they nevertheless seem to include all sorts of phenomena that are representational but nevertheless *non-psychological*.

Consider, for example, circadian clocks: Oscillatory biochemical networks found within the cells of virtually all terrestrial organisms, which have an endogenous rhythm of roughly 24 h that is entrained to the phase of the Earth's day-night cycle by various cues such as light and heat (Gardner et al. 2006). Circadian clocks help to regulate an organism's interactions with the day-night cycle by virtue of *tracking* the current phase of the day-night cycle. For example, the circadian clocks in the flowering plant *Lavatera* allow that plant to reorient its leaves overnight so as to face the 'anticipated' direction of the sun in the morning, thereby maximizing photosynthesis (Schwartz and Koller 1986). The clocks in the *Arabidopsis* allow the plant to muster an insecticidal chemical defense at the time of day the plant is most likely to be attacked by herbivorous insects (Goodspeed et al. 2012). Thus circadian clocks in plants seem to satisfy the conditions of tracking theories: They indicate the current phase of the day-night cycle, and have been selected to control behavior in virtue of that fact.⁵

However, I think it is safe to assume for present purposes that *plants don't have minds*.⁶ So although tracking theories were initially intended to defend materialism by showing that the distinctively mentalistic phenomenon of intentionality is not *irreducibly* mentalistic, they seem to end up entailing that intentionality isn't a distinctively mentalistic phenomenon after all.

⁴ Millikan (1989) also emphasizes the importance of attending closely to how a representation is 'consumed', i.e. how it is used "as a representation" (p.285). But she says little about what consumption consists in, and her conception of consumption is so liberal that the 'pulling' of magnetosomes within the intracellular matrix of a bacterium seems to qualify (see p. 290). But one might wonder if Dretske's use condition is similarly liberal; after all, it was Dretske (1986) who introduced magnetosomes into the philosophical literature as a primitive example of misrepresentation. Whatever Dretske *in fact* thought about magnetosomes, I think his mature theory has the resources to exclude them in a way that Millikan's theory does not. The central reason is that magnetosomes cannot occupy distinct indicator states that can be enlisted to play differential roles in behavioral control.

⁵ Morgan (2014) develops this argument in more detail and extends it to so-called 'structural' theories of representation. Bechtel (2011) argues independently, though on similar grounds, that circadian clocks qualify as representations.

⁶ As I discuss here and elsewhere in the text, there is evidence that plants exhibit remarkably rich forms of adaptive plasticity. Some philosophers and biologists have recently taken this to show that plants are minimally cognitive or intelligent (Gagliano 2017; Calvo Garzón and Keijzer 2011; Maher 2017). It is not always clear whether these proponents of 'plant intelligence' claim that plants literally have minds; one might agree with Calvo Garzón and Keijzer (2011) that plants are "cognitive in a minimal, embodied sense" (p. 166), yet deny that plants are literally psychological subjects with mental states. But those who clearly do argue that plants have mental states tend to do so on grounds that phenomena like representation and learning are inherently psychological (e.g. Maher 2017), which begs the present question. So I think I'm entitled to assume that plants don't have minds for present purposes. Whether we're entitled to assume this *generally* is an interesting question for another occasion.

Some tracking theorists might resist this conclusion. While I don't have the space here to evaluate the nuances of all existing tracking theories, I think it's fair to say that the core phenomenon picked out by tracking theories—states that have the function of carrying information about specific environmental entities—is not psychological in nature. This view is now widely accepted in the literature (Burge 2010; Kriegel 2012; McDowell 1994), and more importantly it is endorsed by tracking theorists themselves. For example, Fodor (1990) writes that “a good theory of content might license the literal ascription of (underived) intentionality to thermometers, thermostats, and the like” (p. 130), and according to Millikan (2000), “Brentano equated intentionality with the capacity to bear a real relation to something nonexistent [but he] was surely mistaken, however, in thinking that bearing a relation to something nonexistent marks only the mental” (pp. 83–4).⁷

So, in sum, we saw earlier that several important projects in the philosophy of mind rest on the hope that a theory of underived content will pick out a distinctively psychological phenomenon: They hope that such a theory will demarcate the appropriate level for psychological explanation (Fodor and Pylyshyn 1988), or will distinguish psychological from non-psychological processes (Adams and Aizawa 2008), or will limit the class of phenomenal states (Bourget 2010). But it seems that the only mature, mainstream theories of underived content currently on offer in philosophy—tracking theories—fail to pick out a distinctively psychological phenomenon.

3 Constancies: representation rementalized?

In his (2010) book *Origins of Objectivity*, Tyler Burge sets his face against the tracking orthodoxy and seeks to *rementalize* representation. He thinks that tracking theorists have dementalized representation because they've misunderstood the demands of naturalization; they've assumed that to naturalize representation we must provide a reductive analysis using notions that are not distinctively psychological, like “[i]nformation, correlation, causation, function” (p. 9). Burge denies that such an analysis is needed, and holds on Quinean grounds that the best possible reason for thinking that representational content is part of the natural world is that it plays an indispensable role in successful scientific explanation. He argues that representational content or accuracy conditions do indeed play an indispensable role in perceptual psychology; specifically, in explanations of perceptual constancy.

⁷ Of the most prominent tracking theorists, I think Dretske (1988) has the best grounds for rejecting the charge of dementalization, though still not *good* grounds. Dretske is centrally interested in how representational content can causally explain the behavior of individual organisms. He thus distinguishes between representations that are selected for behavioral control over phylogeny, and representations that are selected within an organism's lifetime. He focusses on the latter kind of 'ontogenetic' representation, and characterizes it in terms of a specific kind of learning process: Discrimination learning. But this is not the only kind of learning whereby indicators are ontogenetically selected for behavioral control. Gagliano et al. (2016) show that plants can be classically conditioned to grow towards a source of wind that was previously paired with a source of light. For this to occur, the plant must presumably contain some internal indicator of wind direction that is selected for controlling the direction of plant growth, in virtue of what it indicates, *within the plant's lifetime*. This mindless tracking seems to qualify as 'ontogenetic' representation in Dretske's sense.

Since seminal work by von Helmholtz (1866 [1924]), psychologists have sought to explain how we are able to perceive stable properties of distal objects despite noisy, ambiguous, and constantly shifting proximal sensation. Why don't you see an approaching person as getting larger as the image she projects on your retina expands? How do you see her as having a *constant* height, the height she actually has? Helmholtz proposed that perceptual constancy is mediated by 'unconscious inferences' made on the basis of 'assumptions' embodied in our perceptual systems about the most likely distal causes of proximal sensation. So you see an approaching person as having a constant height in part because your visual system compensates for the expanding retinal image on the basis of assumptions about how distant the person is from you. When those assumptions are violated, as with the systematic distortion of distance cues in the famous Ames room, illusions of size inconstancy ensue.

This general idea has been developed into a mathematically rigorous theoretical framework that now lies at the core of mainstream perceptual psychology. Call it the *Helmholtzian Framework*.⁸ As Burge interprets it, the Helmholtzian Framework makes systematic explanatory use of accuracy conditions. It characterizes perceptual constancy as a matter of accurately representing objective features of the mind-independent world, and explains constancy in terms of assumptions that help to fix the conditions under which a perceptual state is accurate. When those assumptions are violated, perceptual states will tend to be inaccurate. Thus, Burge argues, ascribing accuracy conditions to perceptual states isn't just a convenient manner of speaking; it plays a distinctive, systematic and indispensable role in psychological explanation. On these grounds Burge claims that perceptual states comprise natural kinds in psychology that are taxonomized by their accuracy conditions; they comprise a distinctively *psychological* kind of representation.

Indeed, Burge holds that perceptual states are the most *fundamental* kind of representation, in three interrelated senses. First, they are *explanatorily* fundamental in that they play an irreducible role in psychological explanation that cannot be analyzed in terms of notions like information or biological function.⁹ Burge thinks that the informational states picked out by tracking theories do not genuinely represent the world because any putative accuracy conditions they have are explanatorily gratuitous. Second, perceptual states are *psychologically* fundamental since they suffice for an organism to represent objective features of the world. This is important for Burge because he sees it as undermining a long tradition of *intellectualism*, which holds that a subject's capacity for objective representation depends on her possession of sophisticated representational capacities like concepts of self or space. Burge thinks that the conditions of objective representation are more demanding than tracking theorists suppose, and exclude plants and magnetosomes, but are less demanding than intellectualists suppose, and are satisfied by most animals including insects. This relates to a third, *phylogenetic* sense in which perceptual states are said to be fundamental. Burge holds that the evolution of perceptual constancies marks the transition from sensory systems that are merely capable of registering proximal stimuli to systems

⁸ For a few influential expressions of this framework, see Gregory (1980), Knill and Richards (1996), and Palmer (1999).

⁹ I am here using Burge's convention of underlining to indicate the names of concepts.

that can represent distal features of the objective world, and that the appearance of such objectifying representational capacities marks the origins of a primitive kind of mentality. In sum, Burge writes that “perception constitutes, phylogenetically, the first important kind that is representational in a specific, interesting, psychologically distinctive, and probably irreducible sense” (p. 549). As he puts it in a later précis of his book: “Perceptual capacities constitute the most primitive sort of representational mind” (Burge 2014, p. 400).

So Burge (2010) offers an account of the most basic kind of underived representation—perceptual objectivity—which is supposed to mark the origins of a primitive kind of mentality. The account is expressly intended to improve upon tracking theories by *rementalizing* representation. For Burge, states with underived content or accuracy conditions are *ipso facto* psychological. Burge eschews the reductive analyses provided by tracking theorists, so his theory of representation is not set out as an explicit set of conditions. Rather, it takes the form of a detailed characterization of what constancies are and how they work within the Helmholtzian Framework.

Three aspects of Burge’s interpretation of the Helmholtzian Framework are central. First, consider Burge’s characterization of perceptual constancies themselves. Burge (2010) writes that constancies are “capacities systematically to represent a particular or an attribute as the same despite significant variations in registration of proximal stimulation” (p. 408). For example, color constancy “is the capacity to see something as of the same shade of color under very different [illumination conditions]” (p. 387). This talk of ‘seeing’ might suggest that Burge thinks of constancies as capacities for having perceptual experiences with a specific phenomenal character, but he in fact denies that constancies have anything essentially to do with conscious experience and holds that constancies might well mediate the formation of *unconscious* perceptual states (e.g. p. 368).

So what makes constancies *perceptual* if not their essential connection to perceptual experience? What distinguishes constancies from the capacities to track distal entities emphasized by tracking theorists? Burge recognizes that many sensory systems are capable of tracking stimuli by averaging or filtering noisy sensory information. His central example is the salmon’s ability to return to its spawning grounds by following olfactory trails; this is made possible by sensors that track average concentration gradients of specific chemicals. But Burge denies that such processes involve genuine constancies. Constancy involves more than merely averaging or filtering signals that are already present within the proximal stimulus. For Burge, the nature of constancies is tied to the Helmholtzian idea that perception is an *inferential* process that supplements impoverished sensory information. From a contemporary perspective, constancies are mediated by domain-specific computations in sensory systems that systematically extract invariant distal features from noisy and ambiguous proximal stimuli on the basis of prior ‘assumptions’ embodied in the functional architecture of the system.¹⁰ In this way, claims Burge, distal features are ‘separated’ from proximal stimulation, thereby producing states can be accurate with respect to an objective subject-matter.

¹⁰ Burge (2010) is quite explicit that the distinction between constancy capacities and mere sensory registration “hinges on the nature of the internal transformations” (p. 424). See also pp. 408–410 and 424.

This brings us to the second important aspect of the Helmholtzian Framework, the idea that perceptual constancy is mediated by ‘inferences’ and ‘assumptions’ about the world. I use scare-quotes here since there is controversy over the cogency of applying these paradigmatically psychological predicates to sub-psychological systems (Hatfield 2002). For Burge, the ‘assumptions’ underlying constancies are not themselves representational states, but are simply ways of characterizing the domain-specific information tacitly embodied in the architecture of sensory systems that allows those systems to resolve sensory ambiguity. Assumptions merely describe how computation within a sensory system is biased toward the formation of perceptual states that correspond to the most likely distal causes of proximal stimulation. In this way they play a crucial theoretical role for Burge in helping to fix the conditions under which a perceptual state is accurate.¹¹ But assumptions shouldn’t be reified; Burge calls them *formation principles* to avoid the suggestion that they are themselves representational or psychological states. Burge is committed to this non-representational interpretation of assumptions given his view that perceptual states are fundamental in the sense that they don’t depend on the possession of more basic representational capacities.

This leads to the third aspect of the Helmholtzian Framework. We noted that assumptions or formation principles bias sensory processing toward the most likely distal causes of sensory stimulation. We should add: Most likely *for a particular species given the specific environment it evolved in*. The prior structure embodied in sensory systems that resolves sensory ambiguity is generally held within the Helmholtzian Framework to reflect the structure of the environment that those systems are adapted to (Shepard 2001). Burge calls this an organism’s *functional background*: features of the world that an organism must be appropriately related to for the organism to satisfy its biological needs. Importantly, Burge holds that formation principles are specified *in terms of* an organism’s functional background. In this way, an organism’s functional background constrains how formation principles fix the accuracy conditions of perceptual states.¹²

So Burge offers a substantive characterization of what perceptual constancies are, and how they lead to the formation of states with determinate accuracy conditions. On this basis, he holds that perceptual constancies are “necessary and sufficient for [a] system’s being a perceptual system. Their presence is certainly sufficient for perception and objectivity, at least given [an organism’s] functional background” (p. 413).¹³ Burge argues, further, that the explanatory power of ascribing accuracy conditions in perceptual psychology licenses us to hold that perceptual states comprise a distinc-

¹¹ See Burge (2010), p. 285.

¹² See, especially, Burge (2010) pp. 319–324 and 369–370. Burge thinks that this idea dissolves the infamous ‘disjunction problem’ that bedeviled tracking theories. His proposal goes something like this: A rabbit perceives a snake *as* a snake rather than as undetached snake-parts because its sensory systems are biased to track the specific, biologically salient macro-scale features of the rabbit’s environment that help to explain the rabbit’s adaptive success. I confess that I don’t see how this proposal scratches the disjunction problem. Why doesn’t the rabbit see the snake *as*, say, a *death-bringer*? Construed extensionally, that is just as much an adaptively salient macroscopic feature of the rabbit’s environment as a snake. While this specific example might be challenged, the general point is hopefully clear: As Fodor (1990) pointed out long ago, the macroscopic features of an organism’s environment fall under indefinitely many kinds, which might equally well explain an organism’s adaptive success.

¹³ See also Burge (2010) pp. xi, 408, 413, and Burge (2014) p. 399.

tively psychological kind of representation, one that marks the origins of a primitive kind of mentality. Taking these ideas together, Burge claims that bacteria, amoebae and molluscs do not represent the world in a psychologically distinctive way, hence do not have primitive minds, on grounds that their sensory systems do not implement constancies: “They simply react to surface stimulation that is... correlated with environmental attributes” (p. 325). Conversely, Burge (2014) claims that since the oldest sensory systems that do implement constancies are found in certain insects, “these animals have the most primitive type of representational mind. Representational mind begins with bees, spiders, locusts, and preying mantises” (p. 400).

I think Burge offers the most detailed, richly argued, and empirically informed account of a distinctively psychological kind of representation currently on offer in contemporary philosophy. It offers the best current hope for the Content View and the various philosophical projects that depend on it. Over the next two sections I will argue that Burge’s theory in fact dementalizes representation for the same reason tracking theories do. The observant reader might have already seen why. Burge distinguishes psychological representation from mindless tracking by appealing to the explanatory robustness of accuracy conditions. But it was the central task of Sect. 2 to argue that any viable tracking theory must also appeal to the explanatory robustness of accuracy conditions. This helps tracking theories to exclude non-representational phenomena like magnetosomes or salivation, but it doesn’t help them to exclude representational yet *non-psychological* phenomena like circadian clocks in plants. Explanatory accuracy cuts between genuine representation and mere information, not between *mental* representation and mere information.

Before developing this objection, though, I should address one last expository point. *Origins of Objectivity* contains some passages that might appear to be in tension with the interpretation I’ve just provided. Perhaps most notably, just after the aforementioned passage in which he claims that constancies suffice for perception, Burge (2010) writes: “Since they are not characterized independently of the notions representation and perception, one cannot use the notion perceptual constancy as an independent ‘criterion’ to determine when one has a case of genuine perceptual representation... Empirical theory must draw the distinction and identify cases of perceptual constancy” (p.413). Some might take passages like this to show that Burge doesn’t really intend to offer a non-circular account of primitive representation that serves as a criterion for the presence of mentality, and hence that his account isn’t susceptible to the cases of mindless representation I’ll go on to discuss.¹⁴

There’s a benign sense in which this is correct. Burge clearly doesn’t intend to give a noncircular reductive analysis of perceptual representation of the kind championed by tracking theorists. He thinks that such analyses inevitably lead towards dementalization. As I’ve explained, his theory of representation is quite different; what qualifies as a perceptual constancy for Burge is determined by the empirical applicability of the Helmholtzian Framework. Perhaps this makes Burge’s account *virtuously* circular, but it doesn’t insulate his account from counterexamples. Burge is clear from the first page of his book that he seeks to develop an account of what perceptual representation

¹⁴ Thanks to two anonymous referees for highlighting the importance of addressing the interpretive issues raised by this passage and others like it.

consists in that will help to illuminate the origins of a primitive kind of mentality.¹⁵ While he thinks that science speaks to the constitutive questions he's interested in, he recognizes that science doesn't wear its ontological commitments on its sleeve.¹⁶ He thus provides a substantive *interpretation* of the Helmholtzian Framework, which he uses to reach general conclusions about the nature, origins, and scope of mentality: That, say, insects have a representational perspective on the world, whereas bacteria, plants and molluscs do not. Burge makes this claim not on the basis of vague intuition, but on the basis of a substantive theory of what constancies *are* and of how they fix accuracy conditions.

So what does Burge mean when he writes that perceptual constancy is not a *criterion* for the presence of perceptual representation? I think he simply means that he's not providing a reductive analysis of perceptual representation. This is perfectly compatible with his repeated claim that constancies *suffice* for perceptual representation—as it better be, given that he endorses both claims in the same paragraph. And the sufficiency claim is all I need for my counterexamples to have teeth.

In any case, much of this discussion is moot. While I think the textual evidence for my interpretation of Burge is decisive, suppose someone *could* defend an alternative interpretation on which Burge is simply silent on the question of whether mindless systems like cameras or plants have states with underived representational content. On this interpretation, Burge's account would clearly offer no support for the Content View, the view that all underived representation is psychological. However, I'm only interested in Burge's account insofar as it at least purports to vindicate the Content View. So for present purposes, I can simply reformulate my thesis as a conditional: If Burge offers a substantive and fully general account of the most primitive kind of underived representation, then his account fails to pick out anything psychological in nature. I'll assume that the antecedent is true. If someone were to establish that it's false, this would simply provide a faster route to my conclusion.

4 Objection 1: deflating constancies

I now turn from exposition to critique. In this section I present the first horn of a dilemma. I argue that constancies as Burge understands them do not in fact suffice for a distinctively psychological kind of representation. Constancy capacities are found in all sorts of mindless systems. This is not damning for Burge, since he can allow that constancies are at best *diagnostic* of psychological representation. His considered view seems to be that perceptual states qualify as psychological representations not in virtue of being formed through the exercise of constancy capacities *per se*, but in virtue of having explanatorily robust accuracy conditions. But this admission exposes Burge to the second horn of the dilemma, which I discuss in Sect. 5. In short, as we saw in our discussion of tracking theories, all sorts of mindless systems have

¹⁵ Burge (2010) writes that his “aim is to understand the nature of representational mind at its lower border. A corollary of this primary aim is to explain the extreme primitiveness of conditions necessary and sufficient for [perceptual] representation” (p. xi).

¹⁶ See esp. Burge (2010) pp. xv–xvi.

states with explanatorily robust accuracy conditions. Burge's theory thus dematerializes representation for the same reason tracking theories do.

Let's begin by considering the first horn. *Contra* Burge, constancy capacities do not suffice for a distinctively psychological kind of representation. Some form of color constancy processing is now ubiquitous in modern digital cameras, where it is often called *white balance*. Discussion of how to implement color constancy algorithms in cameras and machine vision systems is the stuff of textbooks. For example, Ebner (2007) provides a mathematical introduction to computational theories of color constancy, and surveys the leading algorithms for implementing color constancies machines. Much like Burge, Ebner characterizes color constancy as the capacity "to recognize the color of objects irrespective of the light used to illuminate the objects" (p. xiii). As Ebner points out, the most powerful and widely-used color constancy algorithms are drawn directly from Helmholtzian theorizing in perceptual psychology. For example, Land and McCann's (1971) Retinex theory, which is cited by Burge (2010, p. 351) as an exemplar of the Helmholtzian Framework, has proven especially influential in the field of machine vision.

So color constancies in machines are expressly engineered to operate in accordance with Helmholtzian principles. Specifically, such constancies are mediated by computational processes that systematically extract the surface reflectance of distal objects from proximal sensory signals in which reflectance and ambient illumination are confounded, on the basis of prior 'assumptions' about the spectral properties of the environment. One such assumption employed by many color constancy algorithms is that the world is, on average, gray (Ebner 2007, p. 106). This is no mere 'averaging or filtering', as Burge describes sensory registration; color constancies in cameras and machine vision systems seem to qualify as genuine *perceptual* constancies by Burge's lights. But presumably these machines aren't minded perceivers any more than plants or bacteria are.¹⁷

4.1 Reply 1: cameras don't have constancies

There are two general replies we might consider on Burge's behalf. The first is to deny that the color sensitivities of modern digital cameras qualify as genuine *perceptual* constancies. I've just argued that the sensitivities of cameras *do* qualify as constancies according to the theory of constancies latent in Burge's characterization of the Helmholtzian Framework. What reason might Burge have to deny this?

One such reason might be sought in the idea that constancies involve robust 'assumptions' about the structure of a perceiver's environment. It might indeed seem odd to say that cameras make assumptions. But recall that for Burge, assumptions within the Helmholtzian Framework are not to be understood as explicit representational states. Talk of 'assumptions' is just a colorful way of referring to the prior

¹⁷ Ganson et al. (2012) also argue that constancies as Burge characterizes them do not suffice for genuine perceptual representation, though they don't really explore the implications for Burge's view about the origins of representational mind. Their primary critique points in the other direction: They argue, in a nutshell, that Burgean accuracy-conferring constancies are not *necessary* for perception, and hence that Burge offers no reason to think that perceptual states are essentially contentful. This is distinct from, but compatible with, the argumentative thrust of this paper.

structure embodied within a sensory system that allows the system to systematically extract distal features of the world from noisy and ambiguous proximal stimuli. The constancy algorithms in cameras are clearly mediated by assumptions in this minimal sense—recall the ‘gray world’ assumption mentioned earlier.

A variant of the present reply is to insist that although constancies in cameras are mediated by assumptions, they’re not mediated by sufficiently rich or complex assumptions. There is indeed reason to think that even the most biologically plausible color constancy algorithms do not currently match the perceptual performance of humans or other animals. No current machine vision systems exhibit human-level color constancy (Hurlbert and Wolf 2002). But Burge (2010), like Ebner (2007), endorses the widespread assumption that color perception in animals is computational in nature.¹⁸ So he’s presumably committed to the view that human color perception involves *some algorithm or other*, and offers no reason to think that this algorithm couldn’t in principle be implemented in a mindless machine. Burge offers no reason to think that constancy processes must meet some threshold of integration of complexity to qualify as perceptual; on the contrary, he repeatedly suggests that constancies per se suffice for perception.

This leads to a slightly different way of developing the present reply. Recall that on Burge’s view, constancies suffice for perception *given an organism’s functional background*. The notion of an organism’s functional background plays a crucial theoretical role for Burge in helping to fix the content of the organism’s perceptual states. This suggests that Burge has the resources to deny that cameras and other mindless artifacts are psychological systems on grounds that they are not organisms with a functional background. But there are two serious problems with this proposal. The first is that it renders Burge’s account of primitive mentality unreasonably chauvinistic. It would exclude by definitional fiat what seems like a live empirical possibility: That sufficiently complex artificial systems might in principle be capable of perceiving the world. Burge (2010) in fact seems to reject this kind of chauvinism; he expresses optimism that his “conception of representation can be extended” to artificial systems (p. 293).

However, the second and more important problem with the present proposal is that it mischaracterizes the theoretical role of the notion of functional background within the Helmholtzian Framework. The explanatory point of appealing to an organism’s functional background has nothing to do with a commitment to the dubious metaphysical thesis that only biological organisms can perceive the world. Rather, as Burge emphasizes, the point is to explain how the structure embodied within a sensory system allows the system to track distal features of the world more or less *accurately*. The central idea, recall, is that the structure of the sensory system systematically reflects the structure of the world. In biological organisms this harmony is established through natural selection, but there’s no reason internal to the Helmholtzian Framework to suppose that this is *essential* to constancies. The color constancy algorithms in cameras and other machine vision systems are *engineered* to reflect the spectral properties of the world around us. This is just another way in which the Helmholtzian Framework

¹⁸ He writes that assumptions or formation principles are implemented by “effective procedures, procedures that follow an algorithm” (p. 95).

gains explanatory purchase on reality. So in this sense—the only sense pertinent to the explanatory goals of the Helmholtzian Framework—cameras with color constancy algorithms *do* have a functional background.

4.2 Reply 2: camera constancies don't have accuracy conditions

So Burge seems to lack the theoretical resources to exclude cameras from the scope of his view of constancies. The explanatory power of the Helmholtzian Framework extends further than he thinks. In light of this, Burge might concede that cameras have genuine constancy capacities, but deny that these constancies mediate the formation of states with underived accuracy conditions. That is, he might grant that constancies allow cameras to represent colors accurately, but insist that any representation here is for *us*, not for the benefit of *the camera*: The content of camera states is derivative. Like tracking theorists, Burge holds that for a state to have underived content, its content must help explain what it does. The content of camera states is surely not explanatorily robust in this sense. So Burge would presumably want to deny that the constancies in cameras mediate the formation of states with underived content.

But that would make trouble for his theory of perceptual content. One of Burge's central goals is to establish that perceptual states have representational content *essentially*.¹⁹ As we've seen, Burge argues that perceptual states are individuated by their content on grounds that accuracy conditions play an indispensable role in perceptual psychology. We've also seen that Burge offers a principled explanation of why perceptual states have determinate accuracy conditions: Roughly, sensory systems that tacitly embody domain-specific assumptions about the causes of sensory signals can accurately track those distal causes even when they're confounded in the sensory signal. The natural reading of this proposal is that perceptual states have accuracy conditions *in virtue* of being formed through the exercise of constancies. This is presumably why Burge repeatedly claims that constancies *suffice* for states with underived accuracy conditions and psychological representation.²⁰ But if Burge were to deny that the constancies in cameras suffice for the formation of states with accuracy conditions, he would have no principled account of why constancies in perceptual systems suffice for states with accuracy conditions. His view would be that constancies suffice for perceptual content, *except when they don't*, with no explanation of when or why. This evidently won't do much to support the view that perceptual states are essentially contentful.

Still, the view that cameras might in principle have constancies and yet lack states with genuine underived content seems intuitively correct. It seems hard to deny that cameras and other machine vision systems have constancies as Burge understands them. As we discussed earlier, these systems are expressly engineered to operate in accordance with our best Helmholtzian theories of biological perception. Yet the content of camera states surely derives from our interpretation of them. Ascribing content to the states of digital cameras doesn't help to illuminate the role of those states within the camera. So Burge should abandon the view that constancies suffice

¹⁹ This is the primary target of Ganson et al. (2012). See note 17.

²⁰ See note 13 for references.

for psychological representation. His more basic commitment seems to be to the idea that it is the *explanatory robustness* of a state's accuracy conditions that suffices for the state to be representational in a distinctively psychological sense.

5 Objection 2: mindless accuracy

But why think that the explanatory robustness of accuracy conditions has anything to do with mentality? We saw that tracking theorists also appeal to the explanatory robustness of accuracy conditions to pick out a kind of underived representation, yet by their own admission such representations are to be found in all sorts of mindless systems.

Burge holds that tracking theories fail to pick out a psychological kind of representation because they are expressed using non-psychological notions like information, correlation, or function. However, the notion of accuracy is surely non-psychological too: It applies intelligibly, literally, and indeed paradigmatically to mindless artifacts like watches, thermometers, and fuel gauges.²¹ Indeed, the notion of accuracy seems semantically related to notions like function and correlation. If I went to get my watch repaired and were somehow prevented from using the term 'inaccurate', I might express much the same concept by saying that the hands of my watch no longer *function to covary* with the day-night cycle.

One might suspect that when Burge writes of accuracy, he's employing a proprietary *theoretical* notion, one that does pick out a distinctively psychological phenomenon. But Burge provides no theory that would infuse such an idiosyncratic use of 'accuracy' with meaning. When he writes of accuracy, he simply seems to mean ordinary *accuracy*.²² What is centrally important for Burge is that accuracy conditions be *objective*: Accuracy with respect to some specific feature of objective reality. But of course that doesn't distinguish an idiosyncratic psychological notion of accuracy from plain old accuracy. Watches and thermometers can be accurate with respect to objective reality—that's why we make them.

Burge's claim that states with underived accuracy conditions are *ipso facto* psychological is not based on some special notion of psychological accuracy, but on the claim that ordinary accuracy conditions play an explanatory role in psychology that they don't play anywhere else in science. In particular, Burge (2014) denies that accuracy conditions play a role in biology, writing that "no science takes [a] plant's internal states to have accuracy or truth conditions" (p. 393).

I'll now argue that this claim is false. Accuracy conditions *do* play an indispensable role in biology. We saw in Sect. 3 that biologists explain various activities of plants by appealing to their circadian clocks. Of course it's not just the *possession* of a clock that explains how a plant is able to reorient its leaves overnight so as to maximize

²¹ The cogency of ascribing accuracy to perceptual states is debated (Brewer 2006; Travis 2004), but the cogency of ascribing accuracy to watches and the like is surely uncontroversial. One might point out the accuracy conditions of watches and similar artifacts derive from the mental states of agents. I agree, but this is irrelevant to the point that the ordinary notion of accuracy applies intelligibly to mindless systems.

²² Burge (2010) writes of a "*normal* notion of veridicality... evident in the explanatory practice of perceptual psychology" (p. 308, my emphasis).

photosynthesis, or to release insecticidal chemicals at the time of day when insects are most likely to strike. It's the fact that the clock can be *accurate* with respect to the current phase of the day-night cycle. If a plant's clock were not accurate, the plant would likely shrivel and be eaten.

Plant chronobiologists routinely investigate such counterfactuals in their laboratories by intervening on the accuracy of plants clocks and exploring the consequences—by, for example, raising plants under artificial day-night cycles or genetically manipulating their clocks (Gardner et al. 2006). That accuracy conditions play a central role in explanations of plant activities is not just reflected in the practice of plant chronobiologists, but also in what they explicitly say. For example, Gould et al. (2009) claim that “having a robust and accurate clock increases photosynthesis and productivity in *Arabidopsis*” (p. 899), and Más (2005) describes various genes that play a role in “maintaining the accuracy of [circadian] rhythms” (p. 493). So, following Burge's unimpeachable advice to take the commitments of successful science ontologically seriously, we should hold that biology really does take the internal states of plants to have accuracy conditions.

One objection we can dispense with quickly is that plant circadian clocks do not mediate constancy capacities, so they do not satisfy the conditions of Burge's theory, so they could hardly provide a counterexample to that theory. Burge (2010) indeed argues that circadian clocks are non-representational on grounds that they lack the computational sophistication of genuine constancies. He writes: “The sensitivity to temporal phases can be a product purely of basic rhythms in the organism's body” (p. 520). But this betrays a superficial understanding of how circadian clocks work. Clocks don't keep track of the day-night cycle simply by averaging or filtering a single sensory cue. Information about the current phase of the day-night cycle isn't carried by a single unambiguous signal. Instead, the current phase must be *computed* by integrating several ambiguous cues to entrain the endogenous rhythm of the clock (Gardner et al. 2006).²³ This endogenous rhythm can be thought of as *prior structure* that reflects the diurnal structure of the external world, which allows the clock to extract the actual phase of the day-night cycle from ambiguous sensory input.

In any case, these points are moot. I argued in the last section that Burge should *deny* that constancies suffice for psychological representation. His central reason for holding that states with underived accuracy conditions are *ipso facto* psychological has little to do with constancies, but derives from his view that accuracy conditions only play an explanatory role in psychology. So even if Burge could show that plant circadian clocks do not mediate Helmholtzian constancies, he'd still have to engage with the present reasons for thinking that circadian clocks play an indispensable role in plant chronobiology.

²³ This notion of computing a distal feature from ambiguous sensory information is central to Burge's characterization of constancies (Burge 2010, p. 352). One might worry about applying the notion of computation to plants, whose states presumably don't have the syntactic structure widely assumed to be essential to computation. However, Burge rejects a narrow syntactic construal of computation in favor of the more liberal notion of implementing a computable function (*ibid.*, p. 95). All sorts of systems, including plants, might compute in this more liberal sense (MacLennan 2004), so appealing to computation won't help Burge here.

5.1 Reply 1: circadian accuracy is not explanatory

Burge might be tempted to reply that ascribing accuracy conditions to the states of plants is merely a colorful manner of speaking that might be abandoned without explanatory loss, much like ascriptions of accuracy to magnetosomes or saliva. But there's an obvious difference between plant circadian clocks and, say, magnetosomes: Plant chronobiologists routinely characterize circadian clocks in plants as being accurate or inaccurate with respect to the day-night cycle in the context of rigorous and fruitful scientific explanations, without any indication that they're indulging in a merely colorful, figurative, or otherwise second-rate manner of speaking. None of this is true of magnetosomes or salivation.

Burge's reply here must be supported by reasons for thinking that talk of circadian accuracy is dispensable, reasons that wouldn't also license us to hold that talk of *perceptual* accuracy is dispensable. For that it looks like we'll need a general criterion that would help us distinguish those theoretical postulates that are explanatorily indispensable from those that are not. Here's one conception of explanation on which talk of circadian accuracy would turn out to be dispensable: Assuming that microphysical determinism is true, a Laplacean demon could predict plant behavior on the basis of all the microphysical facts, without ever considering facts about circadian accuracy. But Burge wouldn't want to endorse this criterion, for it would prove too much. If the activities of biological systems could be predicted without ascribing accuracy conditions, then the activities of psychological systems surely could too. Ideal predictability doesn't cut between biology and psychology.

This was surely the wrong place to look for a criterion of indispensability in the first place. Whether or not a theoretical postulate is explanatorily indispensable plausibly turns on whether it affords *understanding* (de Regt 2009). This indeed seems to be Burge's view; he writes that ascribing accuracy conditions to perceptual states is indispensable because it affords a type "of understanding that [is] not attainable apart from attribution of representational notions" (p. 293). But I've argued that ascribing conditions to plant circadian clocks is indispensable on precisely the same grounds: Doing so allows us to understand various capacities of plants we wouldn't be able to understand otherwise, as reflected in the theory and practice of plant chronobiology. Simply denying this would beg the present question. Burge needs a principled *reason* for thinking that talk of accuracy in chronobiology is somehow figurative or second-rate.²⁴

²⁴ Here's a more subtle reply suggested by an anonymous referee. I've characterized Burge's argument that perceptual states have underived accuracy conditions in terms of the indispensability of accuracy conditions as *explanantia* in perceptual psychology. But Burge clearly thinks that accuracy conditions also play a role in specifying the *explananda* of perceptual psychology; he holds that one of the central explanatory goals of perceptual psychology is to explain how animals can accurately perceive specific features of their environments (e.g. Burge 2010, p. 342). This suggests that Burge might reply to the present objection by insisting that the *explananda* of plant chronobiology are not themselves representational in nature. But, again, this begs the present question. Anyone who thought that accuracy conditions play an indispensable role in explaining plant activities would presumably also think that the activities to explained are themselves representational. It *seems* that one of the central goals of plant chronobiology is to explain how plant clocks can accurately represent the day-night cycle. Burge cannot simply *deny* the appearances here—he needs some reason to believe they're illusory.

5.2 Reply 2: circadian accuracy is merely pragmatic

Perhaps such a reason might be found in a distinction Burge makes between *biological* norms, or criteria for evaluating how well a trait contributes to an organism's adaptive success, and *representational* norms, criteria for evaluating the accuracy of a given state regardless of whether it contributes to adaptive success. As Burge (2010) points out, "accuracy is not *in itself* a practical value" (p. 301), and accuracy and adaptive success can sometimes come apart. As the hair-trigger startle response of many birds illustrates, a propensity to produce false alarms can be adaptive when the cost of negligence is death (Godfrey-Smith 1992). So perhaps Burge might grant that normative criteria are indeed indispensable for explaining the activities of plants, but insist that plant chronobiologists mischaracterize these criteria when they describe them in terms of accuracy; they are *biological*, not *representational* norms.

To evaluate this reply we must first observe that biological and representational norms are not mutually exclusive. Burge himself notes that perceptual states can be evaluated with respect to both biological *and* representational norms. The upshot of Burge's distinction is just that biological and representational norms are conceptually distinct, and can in fact come apart. A perceptual state might be accurate *even if* it happens to have maladaptive consequences. So the present reply can't simply be that since circadian clocks contribute to the biological success of plants they *ipso facto* lack genuine accuracy conditions, for that would also license the conclusion that *perceptual* states lack genuine accuracy conditions.

The present reply must instead be something like this: When plant chronobiologists describe certain phases of the day-night cycle as conditions under which circadian clocks would be accurate, they're making a subtle kind of conceptual error. Properly understood, these phases are conditions under which the activities of circadian clocks would eventuate in *adaptive success*. In the case of circadian clocks (and other non-psychological systems) any putative accuracy conditions are in fact just biological success conditions. To paraphrase Burge (2010), ascribing (in)accuracy to plant circadian clocks comes to no more than noting that clocks do or do not serve the organism's needs (p. 410).

Note that this reply entails that plant circadian clocks could not be in states that are (putatively) accurate without those states also being adaptive; on the present view, these come to the same thing. But this is clearly false. A species of herbivorous insect might evolve to take advantage of circadian accuracy in *Arabidopsis* by attacking those plants slightly later than other insects, when the plant's insecticide is normally depleted. In this case, the circadian clocks in *Arabidopsis* would be maximally maladaptive precisely when they're maximally accurate. Accuracy and adaptive success can come apart for circadian clocks just as they can for perceptual states.

What's at issue here, again, is not *whether* plant circadian clocks generally contribute to adaptive success—of course they do, as do perceptual systems. What's at issue is *how* they contribute to adaptive success. Do they play a distinctive kind of functional role, one that is aptly characterized in terms of accuracy? Do they serve to accurately track some specific distal property? Do they function *as* representations for the systems of which they are a part? Plant chronobiology suggests affirmative

answers to these questions. For Burge to assert that plant clocks do *not* contribute to adaptive success by playing this distinctive role is to beg the present question. It is to assert, without justification, that plant chronobiologists are fundamentally confused about how circadian clocks work.

To underscore this point, note that not all circadian clocks contribute to adaptive success by functioning as representations. Consider the humble Somalian cavefish, which has evolved in the perpetual darkness of subterranean caves since their ancestors were cut off from the surface of the Earth roughly two million years ago (Cavallari et al. 2011). Through that process of troglodytism the circadian clocks of cavefish became partly vestigial: They still function to regulate metabolic processes within the fish, but they no longer function to coordinate the fish's interactions with the external world. Relieved of selective pressure to reflect the day-night cycle, they've evolved to have an endogenous period of roughly 47 h.

Here is a case in which the distinction between representational and biological norms would *not* find purchase, a case in which ascribing accuracy to circadian clocks really would be gratuitous. Even if a mutant cavefish were born with a circadian clock that *happened* to have an endogenous period of 24 h, which one day *happened* to coincide with the phase of the day-night cycle, this would be mere coincidence, and wouldn't help to explain any of the fish's activities. If the circadian and diurnal rhythms happened to go out of phase again, the fish's activities would hum along as normal. Not so for terrestrial plants. If the circadian rhythm of *Arabidopsis* failed to coincide with the diurnal rhythm, those plants would be eaten. The explanation of how they avoid being eaten appeals to the fact that the coincidence here is no *mere* coincidence, but a *functioning* co-incidence.

At this point Burge (2010) might complain that by using 'accuracy' to characterize "phenomena that are not, even remotely, distinctively psychological", we'd be using it "so liberally as to debase it" (p. 294). But if 'debase' just means 'applies intelligibly to non-psychological systems', then the ordinary notion of accuracy always was debased—witness watches and fuel gauges. Perhaps Burge has a more rarefied theoretical notion of accuracy in mind that doesn't intelligibly apply to watches or other mindless systems; but then his use of 'accuracy' would be empty, since he hasn't given a theory to infuse it with meaning. He seems to use 'accuracy' in the ordinary sense. The explanatory practice of plant chronobiology provides every reason to think that plant circadian clocks have accuracy conditions in this sense. So by Burge's lights plant clocks should qualify as psychological representations. Yet plants don't have minds. So Burge dematerializes representation in the same way tracking theories do.

6 Diagnosing dematerialization

None of the mainstream theories of underived content on offer in contemporary philosophy pick out a distinctively psychological kind of representation. The Content View seems false, and the various philosophical projects that depend on it seem to be in trouble. What's gone wrong? From a certain perspective, nothing. The dematerializing theories of representation we've discussed articulate clear and theoretically interesting notions of representation, which I suspect are important for understanding the origins

of biological intelligence and the foundations of psychological explanation. They just don't pick out a distinctively *psychological* kind of representation. Like Burge, I think there's an important (though no doubt fuzzy) difference between the tracking states of mindless systems and the intentional states of minded subjects. I just think that he and others have failed to capture it. So there is work to be done to *rementalize* representation. Such a project is beyond the scope of this paper, but in closing I'd like to tentatively diagnose the root causes of dementalization in order to identify a prescription for rementalization going forward.

My diagnosis begins by noting that Burge's theory of representation effectively *just is* a tracking theory. Burge takes himself to be pursuing an entirely different kind of non-reductive project. But it's important to clarify here the sense in which tracking theories are reductive. They don't pursue Nagelian theory reductions (Nagel 1961), nor do they pursue mechanistic reductions (Bechtel 2007b). They are reductive in the sense that they seek to specify conditions under which a state has accuracy conditions using the concepts and theoretical frameworks of natural science, without presupposing the existence of an agent who interprets that state. But Burge's approach is reductive in precisely this sense, too.

Moreover, Burge and tracking theorists endorse the same basic schema for characterizing the nature of representation. They hold that representation fundamentally involves a special kind of causal relation between an internal state of a system and a privileged distal entity in the system's environment. This relation is special in that, unlike ordinary causal relations, the internal state is evaluable as accurate or inaccurate depending on whether it is appropriately caused by the privileged distal entity. Instances of this schema differ in how they spell out the 'specialness' of the representation relation. Burge's theory is distinctive in that it focusses on a specific kind of representation that he takes to be psychologically fundamental, viz. perceptual representation. The theory thus fixes accuracy conditions by appealing to causal relations that are mediated by a specific kind of capacity, perceptual constancy. Tracking theorists are primarily interested in the nature of intentionality per se, so their characterizations of the representation relation tend to abstract away from any specific processes or capacities. For example, Fodor (1990) writes that "all that matters for meaning is [lawful covariation] between symbols and their denotations. In particular, it doesn't matter how that covariation is mediated" (p. 56).

But this is just a difference of emphasis. When tracking theorists focus on specifically *perceptual* representation, they emphasize perceptual constancies in much the way Burge does. For example, Dretske (1981) emphasizes that it's in virtue of constancies that we perceive "*properties of objects*, and not (say) the properties of the retinal stimulation" (p. 163). Fodor (1990) suggests that constancies sometimes *implement* the abstract relation of asymmetric dependence that on his view suffices to fix content (p. 109). Conversely, Burge seems to locate the nature of representation per se in the possession of explanatorily robust accuracy conditions just as tracking theorists do. So I will count Burge as a tracking theorist in what follows.

Why, then, do tracking theories dementalize representation? I suspect they do so because of a subtle conflation of psychological and non-psychological senses of two central concepts, *internality* and *objectivity*. First note that when tracking theorists write of 'internal' representations, they are employing a *mereological* notion of inter-

nality.²⁵ Tracking representations are internal in the sense of being parts or components of a larger system. This seems clearly different from the notion we employ when we talk of the ‘inner’ mental states of a subject. It would presumably be a mistake, for example, to infer from the fact that your watch has internal information-bearing parts to the conclusion that it has inner mental states. We saw in Sect. 2 that sophisticated tracking theories flesh out what it means for an informational state to be ‘internal’ by appealing to the distinctive ways in which those states are enlisted and used within a system such that a robust notion of accuracy finds purchase, but that this doesn’t guarantee that informational states are integrated within the distinctive capacities of *psychological* systems.

The contrast here is helpfully illuminated by Evans (1982) when he writes that for an informational state to be a state of a psychological subject, it must serve

... as the input to a thinking, concept-applying and reasoning system: so that the subject’s thoughts, plans, and deliberations are also systematically dependent on the informational properties of the input. When there is such a link we can say that the person, rather than some part of his or her brain, receives and processes the information. (p. 158)

I think this specific way of framing the requirement is overly intellectualist. I don’t think that an informational state must be integrated within a system’s capacities for reasoning, planning or deliberating to count as psychological. Like Burge, I think that psychological states can be enjoyed by less cognitively sophisticated subjects. Still, I think Evans identifies the right *kind* of requirement. As a rough first pass, we might say that an informational state qualifies as internal in the relevantly mentalistic sense when it is integrated into a unified egocentric frame of reference that allows flexible mappings between the coordinate systems of multiple sensory surfaces and effectors, such that the content of the state systematically coheres with that of other informational states to form global, multimodal states that control the organized, intentional actions of whole organisms. There is considerable neurophysiological evidence that in mammals such unified egocentric reference frames are implemented in the posterior parietal cortex (Cohen and Andersen 2002). If we think of these mechanisms as partial realizers of a subject’s *perspective* or point of view, it arguably starts to become intelligible how the content of an informational state could contribute to the inner perspective of a psychological subject. These ideas are of course speculative and require considerable unpacking, but they are consistent with a good deal of recent work on the mechanisms of perception and action (Andersen and Cui 2009; Morsella 2005; Wu 2011).

This relates to the second of the two notions I mentioned earlier, objectivity. We’ve seen that this notion is important for Burge, and at first blush one might think it would help to distinguish his view from tracking theories, for there is a long tradition of regarding objectivity as a distinctive mark of the mental. The rough idea here is that a system only counts as having a mind when is it capable of representing features of the world *as* objective or mind-independent, for it is only then that the system is one for which a distinction between mind and world finds purchase. In a Kantian strand of this tradition, representational objectivity is explained in terms of a system’s capacity

²⁵ This mereological usage is quite explicit, for example, in Dretske (1988, p. 3).

to grasp an egocentric spatial framework in which entities are represented as standing in specific spatial relations to the the subject (Evans 1982; Peacocke 1992; Strawson 1959). Burge rejects such views as overly intellectualist because he assumes that the capacity to grasp an egocentric spatial framework must be a representational capacity of a subject. He denies that for a subject to qualify as minded it must be capable of *conceptualizing* the conditions for objective representation.

But one might wonder if a shift in the operative sense of ‘objectivity’ has just occurred. We can grasp the traditional notion of objectivity by considering the difference between ordinary object perception and the experience of afterimages. The former state has a kind objective significance that the latter lacks; it presents the world *as* objective and mind-independent.²⁶ Objectivity for Burge simply involves representing features of the world that are *in fact* objective. These seem to be quite distinct phenomena. We’ve seen that circadian clocks in plants plausibly represent objective features of the world (the day-night cycle), but it would presumably be a mistake to infer from this that clocks present those features *as* objective to the plant. So if objectivity simply amounts to representing the objective world, then tracking theorists would have employed a notion of objectivity all along, and objectivity would not serve as a plausible mark of the mental.

So perhaps another way to rementalize representation would be to explore a more restrictive, neo-Kantian notion of objectivity. Burge worries that this will over-intellectualize the mind, but it only does so if one assumes that the only genuine kind of representation is inherently psychological. Once we drop that assumption and recognize genuine *sub-psychological* representation, we are free to hold that a subject might represent features of the world as objective not in virtue of the fact that *she* represents an egocentric spatial framework, but in virtue of the fact that her *nervous system* does.²⁷

A tantalizing possibility is that the two paths toward rementalizing representation that I’ve indicated will converge. Perhaps part of what it is for a state to be internal in the relevantly mentalistic sense is for it to contribute to a subject’s representational *perspective*, and perhaps part of what it is for a subject to have a perspective on the world is for her to apprehend features of the world *as* objective or mind-independent. It is tempting to think that the mutual satisfaction of these conditions would suffice for system with a robust form of mentality, a system with an inner perspective on a putatively outer world. But articulating and defending such a view will have to wait for another time.

Acknowledgements Many of the ideas in this paper percolated under the auspices of the Philosophy of Neuroscience Group at the University of Tübingen. I would like to thank the members of that group, and especially the group leader Hong Yu Wong, for helpful discussion. I have presented this material at Rice University, the University of Graz, and the University of São Paulo, and I am grateful to audiences there for helpful feedback. I am especially grateful to Michael Barkasi, Tyler Burge, Gualtiero Piccinini, Charles Siewert, and two anonymous referees for comments and/or discussion that led to substantial improvements in the text.

²⁶ See Siegel (2006) for rich discussion of objectivity in the relevant sense.

²⁷ For one version of how this idea might be developed, see Grush (2007).

References

- Adams, F., & Aizawa, K. (1994). Fodorian semantics. In S. Stich & T. Warfield (Eds.), *Mental representation: A reader* (pp. 223–242). Oxford: Blackwell.
- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Malden, MA: Blackwell.
- Afraz, S., Kiani, R., & Esteky, H. (2006). Microstimulation of inferotemporal cortex influences face categorization. *Nature*, *442*(7103), 692–695.
- Allen, C., & Hauser, M. (1993). Communication and cognition: Is information the connection? *Philosophy of Science*, *2*(8), 81–91.
- Andersen, R., & Cui, H. (2009). Intention, action planning, and decision making in parietal–frontal circuits. *Neuron*, *63*(5), 568.
- Bechtel, W. (2007a). Biological mechanisms: Organized to maintain autonomy. In F. Boogerd, F. Bruggeman, J. Hofmeyr, & H. Westerhoff (Eds.), *Systems biology: Philosophical foundations* (pp. 269–302). New York, NY: Elsevier.
- Bechtel, W. (2007b). Reducing psychology while maintaining its autonomy via mechanistic explanation. In M. Schouten & D. Jong (Eds.), *The matter of the mind: Philosophical essays on psychology, neuroscience, and reduction* (pp. 172–198). Oxford: Blackwell.
- Bechtel, W. (2011). Representing time of day in circadian clocks. In A. Newell, A. Bartels, & E. Jung (Eds.), *Knowledge and representation*. Palo Alto, CA: CSLI Publications.
- Bermúdez, J. (1995). Nonconceptual content: From perceptual experience to subpersonal computational states. *Mind and Language*, *10*(4), 333–369.
- Blakemore, R., & Frankel, R. (1981). Magnetic navigation in bacteria. *Scientific American*, *245*(6), 58–65.
- Bourget, D. (2010). Consciousness is underived intentionality. *Noûs*, *44*(1), 32–58.
- Brentano, F. (1874 [1995]). *Psychology from an empirical standpoint*. London: Routledge.
- Brewer, B. (2006). Perception and content. *European Journal of Philosophy*, *14*(2), 165–181.
- Burge, T. (2010). *Origins of objectivity*. Oxford: Oxford University Press.
- Burge, T. (2014). Perception: Where mind begins. *Philosophy*, *89*(3), 385–403.
- Calvo Garzón, P., & Keijzer, F. (2011). Plants: Adaptive behavior, root-brains, and minimal cognition. *Adaptive Behavior*, *19*(3), 155–171.
- Cavallari, N., Frigato, E., Vallone, D., Fröhlich, N., Lopez-Olmeda, J., Foà, A., et al. (2011). A blind circadian clock in cavefish reveals that opsins mediate peripheral clock photoreception. *PLoS Biology*, *9*(9), e1001142.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19.
- Cohen, Y., & Andersen, R. (2002). A common reference frame for movement plans in the posterior parietal cortex. *Nature Reviews Neuroscience*, *3*(7), 553–562.
- de Regt, H. (2009). The epistemic value of understanding. *Philosophy of Science*, *76*(5), 585–597.
- Diéguez, A. (2013). Life as a homeostatic property cluster. *Biological Theory*, *7*(2), 180–186.
- Dretske, F. (1981). *Knowledge and the flow of information*. Oxford: Blackwell.
- Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, content, and function* (pp. 157–173). Oxford: Clarendon.
- Dretske, F. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.
- Ebner, M. (2007). *Color constancy*. Sussex: Wiley.
- Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.
- Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.
- Gagliano, M. (2017). The mind of plants: Thinking the unthinkable. *Communicative & Integrative Biology*, *10*(2), e1288333.
- Gagliano, M., Vyazovskiy, V., Borbély, A., Grimonprez, M., & Depczynski, M. (2016). Learning by association in plants. *Scientific Reports*, *6*, 38427.
- Ganson, T., Bronner, B., & Kerr, A. (2012). Burge’s defense of perceptual content. *Philosophy and Phenomenological Research*, *88*(3), 556–573.
- Gardner, M., Hubbard, K., Hotta, C., Dodd, A., & Webb, A. (2006). How plants tell the time. *Biochemical Journal*, *397*(1), 15–24.
- Godfrey-Smith, P. (1992). Indication and adaptation. *Synthese*, *92*(2), 283–312.

- Goodspeed, D., Chehab, E., Min-Venditti, A., Braam, J., & Covington, M. (2012). Arabidopsis synchronizes jasmonate-mediated defense with insect circadian behavior. *Proceedings of the National Academy of Sciences*, *109*(12), 4674–4677.
- Gould, P., Diaz, P., Hogben, C., Kusakina, J., Salem, R., Hartwell, J., et al. (2009). Delayed fluorescence as a universal tool for the measurement of circadian rhythms in higher plants. *The Plant Journal*, *58*(5), 893–901.
- Gregory, R. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *290*(1038), 181–197.
- Grush, R. (2007). Skill theory v2. 0: Dispositions, emulation, and spatial perception. *Synthese*, *159*(3), 389–416.
- Hatfield, G. (2002). Perception as unconscious inference. In D. Heyer & R. Mausfeld (Eds.), *Perception and the Physical World* (pp. 113–143). New York: Wiley.
- Hurlbert, A., & Wolf, C. (2002). Contribution of local and global cone-contrasts to color appearance: A retinex-like model. *Human Vision and Electronic Imaging VII*, 4662, 286–298.
- Jacob, P. (1997). *What minds can do: Intentionality in a non-intentional world*. Cambridge: Cambridge University Press.
- Knill, D., & Richards, W. (Eds.). (1996). *Perception as Bayesian inference*. Cambridge: Cambridge University Press.
- Kriegel, U. (2002). Panic theory and the prospects for a representational theory of phenomenal consciousness. *Philosophical Psychology*, *15*(1), 55–64.
- Kriegel, U. (2012). Personal-level representation. *Protosociology*, *28*, 77–114.
- Land, E., & McCann, J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, *61*(1), 1–11.
- Maclennan, B. (2004). Natural computation and non-Turing models of computation. *Theoretical Computer Science*, *317*(1–3), 115–145.
- Maher, C. (2017). *Plant minds: A philosophical defense*. New York, NY: Routledge.
- Más, P. (2005). Circadian clock signaling in arabidopsis thaliana: From gene expression to physiology and development. *International Journal of Developmental Biology*, *49*(5–6), 491.
- McDowell, J. (1994). The content of perceptual experience. *Philosophical Quarterly*, *44*(175), 190–205.
- Mendelovici, A. (2013). Reliable misrepresentation and tracking theories of mental representation. *Philosophical Studies*, *165*(2), 421–443.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Millikan, R. (1989). Biosemantics. *The Journal of Philosophy*, *86*(6), 281–297.
- Millikan, R. (2000). Naturalizing intentionality. In B. Elevantch (Ed.) *The proceedings of the twentieth world congress of philosophy*, vol. 9, (pp. 83–90). Philosophy Documentation Center.
- Morgan, A. (2014). Representations gone mental. *Synthese*, *191*(2), 213–244.
- Morsella, E. (2005). The function of phenomenal states: Supramodular interaction theory. *Psychological Review*, *112*(4), 1000–1021.
- Nagel, E. (1961). *The structure of science: Problems in the logic of scientific explanation*. New York, NY: Harcourt.
- Palmer, S. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
- Peacocke, C. (1992). Scenarios, concepts and perception. In T. Crane (Ed.), *The contents of experience: Essays on perception* (pp. 105–135). Cambridge: Cambridge University Press.
- Ramsey, W. (2007). *Representation reconsidered*. Cambridge: Cambridge University Press.
- Rorty, R. (1979). *Philosophy and the mirror of nature*. Oxford: Blackwell.
- Schwartz, A., & Koller, D. (1986). Diurnal phototropism in solar tracking leaves of *lavatera cretica*. *Plant Physiology*, *80*(3), 778–781.
- Shea, N. (2007). Consumers need information: Supplementing teleosemantics with an input condition. *Philosophy and Phenomenological Research*, *75*(2), 404–435.
- Shepard, N. (2001). Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences*, *24*(4), 581–601.
- Siegel, S. (2006). Subject and object in the contents of visual experience. *Philosophical Review*, *115*(3), 355–388.
- Sterelny, K. (1995). Basic minds. *Philosophical Perspectives*, *9*, 251–270.
- Strawson, P. (1959). *Individuals: An essay in descriptive metaphysics*. London: Routledge.
- Travis, C. (2004). The silence of the senses. *Mind*, *113*(449), 57–94.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.

- van Gelder, T. (1995). What might cognition be, if not computation. *The Journal of Philosophy*, 92(7), 345–381.
- von Helmholtz, H. (1866 [1924]). Concerning the perceptions in general. In J. Southall (Ed.) *Treatise on Physiological Optics* (3rd ed.), vol. 3, (pp. 143–172). Rochester, NY: Optical Society of America.
- Wu, W. (2011). Confronting many-many problems: Attention and agentive control. *Noûs*, 45(1), 50–76.