
PERCEPTION AND THE ORIGINS OF TEMPORAL REPRESENTATION

BY

STEVEN GROSS

Abstract: Is temporal representation constitutively necessary for perception? Tyler Burge argues that it is, in part because perception requires a form of memory sufficiently sophisticated as to require temporal representation. I critically discuss Burge's argument, maintaining that it does not succeed. I conclude by reflecting on the consequences for the origins of temporal representation.

1. Introduction

Is temporal representation constitutively necessary for perception? That is, is it part of the nature of perception that a creature's having the capacity for perception requires that it have the capacity to represent some sort of temporal phenomenon – whether temporal order, or duration, or place in a temporal cycle, or particular times? Tyler Burge (2010) argues that it is, in part because perception requires a form of memory sufficiently sophisticated as to require temporal representation.

Temporal representation is pervasive, fundamental, and perplexing. But Burge's intriguing contribution has to date not garnered the attention it deserves. In what follows, I critically discuss Burge's argument. I argue, first, that perception does not require the capacity to carry forward its content in time; and, second, that, even if it did, there can be cases in which no explanatory gain accrues from assigning the state *temporal* content.

If Burge's argument does not succeed, a question arises. The aim of his book is to understand and explain the constitutive origins – that is, the

minimal constitutive conditions – of objective representation, the most basic form of which is perception of the physical world. The argument on which I focus is of interest because it embodies an account of the constitutive origins more specifically of *temporal* representation, at least as it arises in connection with perception. If this argument fails, and if temporal representation is *not* constitutively necessary for perception, then what are its origins? What must minimally be the case for perceptual representational content to include temporal elements? I conclude by suggesting that the nature of temporal representation may not determine particularly informative minimal constitutive conditions.¹

2. *Some Burgean themes*

Before turning to the argument, I quickly review some Burgean themes that frame his discussion. As far as possible, I will be granting this framework, so that the objections raised later will have force even by Burge's own lights. Accordingly, I here simply lay out the framework *sans* defense.

Burge distinguishes among mere functional sensory registration of information, perception, and propositional thought (e.g. pp. 104–5, 315–9, 421–36). Sensory registration of information involves a sensorily-mediated statistical correlation, counterfactual dependence, or causal dependence of a mind-brain state with, or on, something external. But mere sensory registration of information, even if functional, does not suffice for a state's having representational content, according to Burge. So, perception, the most basic form of representation, differs from mere sensory registration of information: only the former possesses genuine correctness, or accuracy, conditions. Burge thus sets himself against those – such as Fodor (1990), Dretske (1988), Millikan (1989), and Gallistel (1990) – who would naturalize representational content by reducing it to the obtaining of correlational, causal, counterfactual, or bio-teleological relations (pp. 292–308).

Perception differs as well from propositional thought. For example, the latter's role in inference underwrites the ascription of propositional structure, while with perception, though we are warranted in ascribing a demonstrative and an attributive aspect, we are not, according to Burge, warranted in ascribing propositional structure.² Burge argues, further, that a capacity for perception – and thus for being in states with representational content – does not require a capacity for propositional thought, far less a capacity for self-conscious, or reflective, propositional thought. So, Burge also sets himself against those who, in his view, over-intellectualize the requirements for representational content by requiring such cognitive sophistication for perception (pp. 208–9, 264–88) – for example, Strawson (1959), Davidson (2001), and Brandom (1994).³

What warrants an ascription of representational content? Burge requires that it be necessary for explanation. While cases of mere sensory registration are indeed often *characterized* in representational terms (and can even allow for talk of misrepresentation, since a correlation, e.g. need be merely statistical), explanation does not *require* any attribution of content in such cases. Our best perceptual science, however, has uncovered a natural joint; for there is a natural psychological kind – perceptual states – our best explanations of which *do* require the attribution of representational content. In particular, our best scientific explanations of how perceptual systems solve inverse problems and maintain perceptual constancies posit representational content. ('Inverse problems' refers to the one-many relationship between what's sensorily transduced and its distal causes. 'Perceptual constancies' refers to capacities to perceptually represent the same distal feature across variously changing ambient conditions.) Burge thus also stands against those who deny that our best theories in the perceptual sciences in fact advert to representational content in their explanations (pp. 296–8) – e.g. Chomsky (1995, 2003), to which Burge (2003) replies.

To clarify, let's briefly consider some examples. Here's a greatly simplified explanation of salmon homing (pp. 425–7): As they swim towards their breeding grounds, their heads move from side to side, which enables them to sample the water on each side for chemical traces left behind during their migration to sea; they then head in the direction with the greatest concentration. There results a correlation between their internal state and the direction of their breeding ground – so, a functional sensory registration of information. Indeed, since their heading is maintained through variation in the ambient environment, we can even say that this mechanism preserves a constancy. But it is not a *perceptual* constancy, according to Burge. For it suffices to explain the sustaining mechanism that one characterize, as I just did, the sampling mechanism and its connection to motor activity (albeit of course in more detail). Positing representational content – as of *that direction* – would be explanatorily otiose. It is otherwise, however, when it comes to explaining, for example, the visual system's ability to maintain lightness constancy across changing ambient light – as when the appearance of newspaper print more-or-less remains the same indoors and outdoors despite orders of magnitude more light being reflected from the black print outdoors than is reflected by the white background indoors. Here, Burge argues, our best scientific explanations posit operations over representations (pp. 351–4).

Indeed, Burge singles out perceptual constancies as providing the mark of perception (pp. 408–13). They are, he argues, sufficient and probably necessary for genuine perceptual representation. The necessity claim is not that perception as of X requires mechanisms sustaining perceptual constancies with respect to X. Sometimes, the necessity is satisfied

derivatively via perceptual constancies for *something else*. Temporal representation is one such case, to which we return shortly. (Burge discusses numerosity in this regard as well.) That perceptual constancies are *sufficient* for perception is in fact trivial since for a constancy to count as *perceptual* its sustaining mechanisms must operate over representations.

This makes it clear, as Burge emphasizes, that perceptual constancies do not provide the basis for an analysis of perception into non-representational terms; nor do they provide a criterion that might settle disputed cases: any characterization of perceptual constancies that would distinguish them from non-perceptual ‘constancies’ (such as that arguably exhibited in salmon homing) must deploy representational vocabulary (pp. 397, 408, 413).⁴ But Burge rejects in any event demands for a non-representational characterization of representational content – in addition to rejecting, as we have seen, particular proposed naturalizations. In his view, representational notions earn whatever naturalistic *bona fides* they require from their role in explanation.⁵

3. *Burge’s argument*

Let’s turn now to Burge’s argument for the claim that a capacity for temporal representation is constitutively necessary for perception. The basic idea is that a condition on perceptions’ having content is that some of them are used, and this use sometimes takes time. Carrying these representations forward in time amounts to a form of memory, one that requires a sensitivity to temporal order (for example, differentiating what one perceives *now* from what one perceived *before*). But when such sensitivity is ‘harness[ed]’ (pp. 522, 527) to a capacity for perception, it itself rises to the level of representation and gets encoded into representational content.

In Burge’s own words, and with his numbering (albeit switched to Roman to avoid confusion with page references), the argument runs as follows:

- (i) A capacity for perception must be associated with a capacity for guiding representational use of perception.
- (ii) [Such] use cannot always be simultaneous with perception. ...
- (iii) For perception to guide use, past perception must be coordinated through memory with present or anticipated representational use.
- (iv) This coordination requires sensitivity to temporal order. ...
- (v) ... [T]he systematic coordination of temporal sensitivity ... with perceptual representation suffices for temporal representation.
- (vi) Therefore, a capacity for perception must be associated with a capacity for temporal representation (p. 527).

Premises (i) and (v) in particular require further elaboration. First, ‘use’ in (i) is intended broadly. In particular, it is not limited to intentional action, but can include such passive behavior as the formation of an emotion, blushing, or *schreck* reactions (the passive shutting down of behavior, as when a deer gets stuck in the headlights) (p. 526). We’ll take up in a moment what it is for such use to be itself representational and to be guided. The rationale behind (i) is that such use is a necessary condition for perceptions to have sufficiently *determinate* representational content (pp. 319–26). The candidate contents of perception, on Burge’s view, are constrained by aspects of an organism’s ethology: its environment and what it does in that environment to satisfy its needs. Only if perceptual representations find such a use can ethology set such constraints, and only then can there be sufficient determinacy for perceptions to have any content at all. Burge does not hold that ethological considerations alone can turn aside or answer *all* questions of content under-determination. For example, he holds that one must appeal to other considerations in order to answer whether a frog’s visual perceptions are as of flies or as of black objects sized, shaped, and moving in such-and-such ways. But he does hold that ethological considerations can eliminate from contention many philosophers’ ‘contrivances,’ such as the perceptual analogues of Quinean candidates for ‘gavagai’ (pp. 323–4).

In a qualified sense (to which we return in Section 6), (v) articulates Burge’s conception of the origin of temporal representation. Recall that mere functional sensory registration of information is necessary but not sufficient for perception. Only with the latter do we have something with genuine accuracy conditions, and we are justified in assigning such conditions only when explanation requires it. Such assignments are indeed required – it is the mark of perception – when there are perceptual constancies. But, for some kinds of representation, the necessity of perceptual constancies is indirect. Such is the case with temporal representation (p. 521). Temporal representation, including temporal representation in perception, does not require perceptual constancies with respect to times or temporal features. There need not be mechanisms that filter changes in other ambient contributions to proximal stimulation in order to fix a representation of a distal temporal feature.⁶ Rather, according to Burge, temporal representation arises when a temporal *sensitivity* (a functional correlation between something temporal and states of the organism)⁷ is ‘harnessed’ to the perception of *something else*. It is in this sense – because it is a perceptual constancy with respect to something else – that the perceptual constancy it requires is indirect.

The particular harnessing to which Burge’s argument appeals is spelled out in premises (ii)–(iv). The use of a perception follows it in time: a perception had *then* is acted upon *now*. The coordination of perception and use thus requires carrying representational content forward in time, a kind of memory: the memory state whose content is that of the

perception had *then* is not the perception that is had *now*. And this coordination, Burge claims, in turn requires, or amounts to, a sensitivity to temporal order.

This last claim receives little development, but perhaps can be defended as follows. Recall that temporal sensitivity requires a functional correlation between temporal features and organismic states. In discussing more specifically sensitivity to temporal order, Burge writes of the ability to 'discriminate and retain an order of stimulations that corresponds to an order of events' (p. 518). This suggests that the organism be capable of being in a state that corresponds (or tends to correspond) to its having been first stimulated in one manner and then in another – and thus tends to correspond as well, at least in the normal basic case, to the order of distal events that caused the stimulations. But an organism capable of perception, and thus of memory states that retain the contents of perceptions, seems to meet this condition, since the stimulation that leads to a memory will proceed the stimulation that causes a perception concurrent with that memory. To spell this out, let S1 be a stimulation that gives rise to a perception P1 with a certain content – similarly, for S2 and P2. Let M1 be the memory state that retains the content of P1. Such an organism will be in the complex state P2 + M1 just in case it receives stimulation S1 before S2 (assuming requisite background conditions are in place: S1 and S2 are suitably temporally proximate, the organism is awake, etc.). In this sense, the complex perceptual-memory state corresponds to an order of stimulations (and thus in the normal case to the order of distal causes), which amounts to a basic kind of sensitivity to temporal order.⁸

If this is right, the capacity for use required for sufficiently determinate representational content in perception itself requires a sensitivity to temporal order. (If it is not right, there may be a problem with Burge's argument beyond those I press in what follows.) In this sense, Burge claims a temporal sensitivity is harnessed to the perception of something *else* and thus, according to premise (v), rises itself to the level of representational content. It's not just that one now perceives [that, red], but one also has a perception of [that, red, now].

It may not be obvious why this harnessing of temporal sensitivities to the perception of something else must or should confer representational status. That is the subject of Section 5 below. But one may view Burge's discussion as proposing a natural joint at which to locate the origins of temporal representation – motivated by discussion of specific cases. The implicit challenge, if one disagrees, is to specify how *else* temporal representation might first arise. We return to this question in Section 6. The remarks above are intended to clarify the premises of Burge's argument and their place in his views, not to defend them (premise (iv) excepted). In what follows, I raise questions especially about (ii) and (v), but also about the others.

4. *First objection: instantaneous use*

According to Burge, an organism capable of perception can be very simple. For example, he holds that a capacity for *spatial* representation, unlike that for temporal representation, is *not* constitutively necessary for perception:

An animal could have a visual system capable only of attributing lightnesses or colors. ... [It] lacks any capacity to attribute any spatial relations, including spatial relations between an occurrence of lightness or color and the animal itself. ... Perhaps the animal closes its shell in response to instances of some colors, or some achromatic spectrums of lightness, and opens its shell in response to others (pp. 496–7).

He argues as well that we lack grounds for holding that perception constitutively requires a capacity for perceiving change (for example, for tracking movement), and thus that we aren't in a position for arguing on *this* basis that perception constitutively requires a capacity for temporal representation:

Perhaps there could be a perceiver capable only of awakening and having a perception of color or shape. If one color or shape is perceptually represented, the perceiver goes back to sleep. If another color or shape is represented, the perceiver opens its mouth and receives nutrients flowing by, and goes back to sleep. Perhaps such an individual or system could have perception but lack a capacity to perceptually track change (p. 528).⁹

My first objection suggests that, if we grant these claims, we should allow similarly simple organisms capable of perception but not temporal representation. (Those not willing to grant Burge's claims can treat what follows as *ad hominem*.) Consider a very simple organism capable of perceiving only occurrences of lightness. Burge's premise (ii) maintains that, no matter the organism's simplicity, the organism's use of its perception cannot always be instantaneous. But why must this be so? It is not obvious what precludes there being an organism whose entering the very state that realizes the perception just is entering a state that is also a functional use of that very perception. For example, entering that state could involve phosphorescence or an electro-magnetic change. The propagation and consequent effects might take time, but the change itself could be coincident with the perception. And any such change could have functional significance. It could initiate movement away from a predator associated with the perceived light or towards a food source associated with it; signal a predator in a way that causes it to abandon the hunt; open an orifice or initiate a molecular process; etc. An organism that instantaneously puts its perception to use would have no need for even the rudimentary kind of memory to which Burge's argument adverts: the retention of perceptual content over time from its formation to its use. These cases thus give us

reason to reject Burge's claim that '[e]ven in cases where perception effects what we are inclined to call "immediate" release of simple inflexible instinctual behavior, there is necessarily at least a brief time lag between the occurrence of the perception and impetus to a use' (p. 527). At least with respect to these organisms, speaking of immediate release would be accurate.

One might reply that such uses would not be *representational* uses, as premise (i) requires. This requirement might lend itself to two readings: as a requirement that the perception's representational aspect be made use of, or as a requirement that the use itself be representational. Burge speaks of the latter (p. 526), but it may be that the two requirements come to the same thing. What makes such a use itself representational might just be that it 'makes use' of the perception's representational content, and 'making use' of the perception's representational content might require – or even, as Burge puts it, 'engender' (p. 547) – representational content in the use. This thought sits naturally with the idea that there is only such content at all if explanation posits law-like transitions from such perceptual contents to such use contents. But now Burge's contention that the use of perceptual content 'engenders' use content suggests that we lack an independent and antecedent way of determining whether the use has representational content. It's thus hard to see how one would *motivate* the claim that, in our cases, the uses are not representational. It certainly is easy enough to *assign* them content: *move thusly, go away, eat, or absorb*. Moreover, Burge's broad construal of use – encompassing emotions, blushing, and *schreck* responses (or their initiation) – suggests a similarly liberal construal of *representational* use.

A different reply to 'instantaneous use' cases objects that such uses aren't *guided* by perception – as, again, is required by premise (i) – regardless of whether the use has representational content. A natural first thought here is that, for any token state, if it is to guide another, then it must precede it in time and not be (wholly) coincident with it. But this conception of guidance does not seem to be Burge's. Premise (ii) maintains that guided representational use cannot *always* be simultaneous with perception – which suggests that *sometimes* it *can*. (Cf. also: 'use cannot both be guided by perception and *always* co-occur with it instantaneously' – p. 527, my italics.) So Burge, at least, cannot appeal to this first thought. Perhaps Burge's view is that guidance is a 'type' phenomenon that requires that *some* token pairs not be coincident, even while allowing that some can be. This would have to be made out: what might lie behind a conception of guidance that requires *all* token pairs to involve a time-lag is perhaps clearer than what might lie behind one that requires a time lag only for some. In any event, deploying here either of these conceptions of guidance requires defense. Arguably, there are less demanding conceptions of guidance that, unlike the two more demanding conceptions sketched above, would not require temporally-extended transitions between representation states at all. Consider computational models that generate states through the

simultaneous satisfaction of variously structured sets of constraints, but where the constraints' structure allows for some notion of dependency.

Moreover, if we deploy one of the more demanding conceptions that require a temporal lag, a problem arises concerning premise (i). (i) says that a capacity for perception must be associated with a capacity for guiding representational use of perception. Its rationale is that only then will there be sufficient determinacy of content. But if the reference to guidance is understood in a demanding sense, then (i) outruns its rationale. For whether there is a time-lag in use is *irrelevant* to whether the use sufficiently connects perception with the organism's ethology to place sufficient constraints on its content. The constraints imposed are the same whether the use is instantaneous or (even just slightly) delayed.¹⁰

Perhaps Burge might allow that the concern for sufficient determinacy provides only a partial motivation for (i) so construed, but suggest that the guidance requirement is independently motivated. The motivation would have to be supplied. But, in any event, the problem would remain that (i) is dubious if guidance is construed as requiring a time lag – whether always or sometimes. To make this vivid, consider the following two species. Let the first be one of the simple organisms that Burge himself describes in the quotes above. There is a time lag in use, so no hackles are raised on that score. Suppose, after much empirical investigation, the organism's perceptual system comes to be very well understood. Its constancies are identified and the computational mechanisms that sustain them, as well as the physiology that realizes those mechanisms, are systematically studied. Moreover, further empirical research provides as well a similarly detailed understanding of its use system(s). Then, the second species is discovered. It appears very similar to the first in many ways. In particular, detailed investigation seems to reveal that it possesses a perceptual system that's a close homologue of the first species'. However, as research moves to its use system, it's discovered that the species is one of my instantaneous users. If we uphold (i), with guidance construed in a demanding way, we must deny that the second species perceives. Its sensory registrations fail to rise to the representational level, and its constancies are not *perceptual* constancies, despite the apparent homology – just because there's not a (say) 100 ms lag in its use.

5. *Second objection: what explanation requires*

So far I have raised a doubt about premise (ii) and have argued that one way of trying to turn aside that doubt renders premise (i) questionable. We might recast this first objection as a dilemma: if guidance doesn't require a temporal lag, we may reject (ii); if it does, we may reject (i). But suppose

you remain unconvinced. There is a further objection, even if use must be guided in a way that entails a time lag and a sensitivity to temporal order. So, suppose we put aside the instantaneous use cases and allow the requisite time lag. Still, one can ask whether such sensitivity suffices for temporal representation. That is, one can question premise (v).

Burge argues for the general claim embodied in (v) by considering some examples mentioned above: an animal perceptually tracking a moving particular and a bee associating a phase in a circadian cycle with a perception of sugar water (pp. 521–2). It is compelling that these organisms have the capacity for temporal representation. But it's unclear that this suffices for the general claim. The worry is not that we need deductive, not inductive, grounds. If the cases were representative of the whole, the argument would be very strong indeed. But it's this that's unclear. Burge characterizes the perception of change or movement as the 'simplest sort of temporal representation in perception' (p. 521). But we saw that he himself argues that such a capacity is not necessary for perception (pp. 528–9). My worry is thus that, by his own lights, there are simpler cases not covered by Burge's discussion, even if we allow a (brief) time lag between perception and use. Why think, in such cases, the temporal sensitivity implicated in perception suffices for temporal representation?¹¹

It is only minimally interesting to note that there's room to worry that Burge hasn't provided sufficient grounds in *favor* of (v). More interesting would be a consideration *against* (v). Indeed, I maintain that there is a *Burgean* consideration against it. Recall that it's a central feature of Burge's conception of representational content that one is warranted in ascribing it only when doing so plays some *explanatory* role. So, it would be an objection to (v) if there are cases where one could explain the relevant psychological functions, including the role of temporal sensitivity in them, without positing temporal representational content. For then positing temporal representational content would be explanatorily *otiose*. And indeed this seems to be the case.

We may suppose that our simple organism is in but one perceptual state at a time and thus that there is but one perceptual state at a time that is poised for use. Let us suppose as well that the time lag between perception and use (and indeed between use and ethologically relevant consequences) is so short as to be negligible so far as ethological success is concerned. For example, perhaps the organism perceives red and as a result opens an orifice that remains open till satiation – or does nothing if already satiated. The redness signals the presence of a nutrient absorbed through the orifice. The organism's environment is such – let's stipulate – that normally the nutrient is present for at least 5 seconds after its associated color is perceivable and in an amount beyond that required for satiation; and the organism is such that satiation is reached in 2 seconds (so, normally before the nutrient is no longer present). Now, suppose the time course from perception to use to

the beginning of absorption is 100 ms. In this case, no explanatory purpose is served by assigning temporal representational content. The time-course of the organism's states, together with the rate of relevant environmental change, suffices to explain the use of the organism's perceptual states and their ethological function. No explanatory power is added by, for instance, including *now* in the perception's representational content and having it change to *then* or *very recently* during the time lag from perception to use. Given the environmental facts, it suffices to explain ethological success that normally it is the present color that the organism perceives and that its use does occur very soon after.

Let me clarify the basis of my otioseness claim by distinguishing it from a different basis one might employ but that is not mine. This other argument suggests that temporal representation in such cases is otiose from the *organism's perspective*, or *for* the organism. The idea is that representational content is a limited resource, because, willy-nilly, greater representational complexity requires a greater expenditure of energy. So, if including temporal representation increases representational complexity, but the organism can achieve its ethological goals without it, then *ceteris paribus* it will not deploy such representations. So, *ceteris paribus* we should not ascribe it. Why should one think that greater representational complexity requires a greater expenditure of energy and that temporal representation increases representational complexity? One way of cashing out these claims that might lend them support would invoke a Language of Thought hypothesis according to which representations have a language-like syntactic structure that is mirrored in its realizing neurophysiological vehicle. Burge, however, rejects – or at least sees no support for – such a hypothesis at least in the case of perceptual representations and probably for any representations had by an organism incapable of propositional thought (pp. 39 fn. 15, 95 fn. 45). He would thus reject at least this basis for pressing an otioseness claim. It's thus important to see that it is not *my* basis. *My* claim is that ascribing temporal representation in such cases is otiose from our *own* explanatory perspective, *regardless* of the organism's own limited resources.¹²

A possible reply to the claim of otioseness is that content is assigned to capture the (wide) functional role of a representation. So, insofar as characterizing the functional role of even our simple organism's representations requires adverting to their temporal features, these temporal features should themselves be reflected in their representational content. But, even if functional role *determines* content, it is far from clear that one should accept the general principle that all aspects of functional role are a *part* of content. Indeed, Burge himself rejects it. Consider his characterization of the functional role of perception. Perception, according to Burge, functions *inter alia* to register what is *causally present* (p. 378). Does it follow that it must represent the causally present as such – in

particular, that we can perceptually attribute *causation*, and that indeed *every* perception represents what it demonstratively picks out as the cause of this very perception? Such a view has been endorsed by Searle (1983). But Burge (1991) argues at length against this position. Of his two main objections, the one that is pertinent to our current question is that it is simply a mistake to think that conditions on a perception must be reflected in perceptual content. To simply claim otherwise would beg the question. In any event, it wouldn't suffice by Burge's own standards, which require that reflecting functional role in representational content yield some explanatory gain. Note further that Burge holds that perception functions to register the 'here and now' (p. 542). Were this feature sufficient to ascribe a *representation* as of *here* and *now* to each perception, the necessity of temporal representation for perception would follow trivially, without reference to a temporal lag in use. But this is not how Burge argues.¹³

Finally, it might be objected that we have misdirected our challenge. One might try retaining premise (v) – that the systematic coordination of temporal sensitivity with perceptual representation suffices for temporal representation – by suggesting that our case in fact does not involve an appropriate 'harnessing' of a temporal sensitivity. If it can be argued that such organisms lack the requisite *sort* of coordinating memory, or that in such cases the systematic ordering of stimulus, perception, retained content, and use does not actually amount to a temporal sensitivity, then one might instead reject (iii) or, despite our earlier attempt to render it plausible, (iv). Of course, this would not save Burge's argument. But whether (v) can be saved matters to the question of temporal representation's origins, to which we are about to turn in our concluding section.

Note first, however, that we have not just objected to Burge's particular argument. By providing counter-examples to its conclusion, we have argued for its negation. If one would advance some other argument for the constitutive necessity of temporal representation for perception, one would need also to diagnosis ours: presumably, an alternative argument for Burge's conclusion, if successful, would entail the impossibility of our simple organisms.¹⁴

6. *Conclusion: the origins of temporal representation*

Having rejected Burge's argument, what should we say about the origins of temporal representation? First, we need to clarify the question. In Section 3, we said that premise (v) articulates Burge's conception of the origin of temporal representation – in a qualified sense. 'Origin' here might refer at least to phylogenetic, ontogenetic, or constitutive origins. Our focus so far has been Burge's argument that it's *constitutive* of perception that a capacity

for it requires as well a capacity for temporal representation. But Burge denies that it's constitutive of temporal representation that it depend on perception. Temporal representation, on his view, requires a capacity for *de re* application of attributives to particulars, but this could be accomplished non-perceptually in pure thought (pp. 523–4) – for example, in temporally ordering one's own thoughts (I was thinking about time, but now am thinking about numbers). But Burge does endorse (v)'s 'origins' claim construed phylo- and ontogenetically of actual animals. Further, one might endorse (v) construed as a constitutive claim concerning temporal representation *in perception*.

So, having rejected Burge's argument, what should we say about the origins of temporal representation, construed in these qualified ways: how does temporal representation arise phylo- and ontogenetically, and what must be the case, as a constitutive matter, for perceptual content to include temporal elements? The phylo- and ontogenetic questions are obviously empirical matters, but they may be constrained by how things stand constitutively. Following Burge, we may of course answer: the ascribed temporal content must play some explanatory role, which it would not with our simple organisms. But this does not tell us much, since constitutive claims *generally*, for Burge, obtain in virtue of 'natures,' which are relatively fundamental *explanatory* kinds (p. 57). What more specifically must be the case for temporal representation in perception to figure in genuine explanation?

If one could retain premise (v) in the manner suggested at the close of Section 5, then one would retain as well a Burgean answer to the origins of temporal representation in perception: it involves the 'harnessing' of a temporal sensitivity to perception. It's just that we now see that not all individuals capable of perception exhibit such harnessing, contrary to premise (iii) and/or premise (iv). The obtaining of these premises, however, along with (ii), are what give content to Burge's talk of harnessing in the first place, at least as it is deployed in the argument's general case. So, absent some (alternative) substantive accounts, it would remain obscure what further conditions carrying forward content in time must satisfy to amount to memory *of the appropriate sort*, and what further conditions a correlation between some temporal features and states of the individual must meet to amount to a *genuine* temporal sensitivity. Indeed, it's unclear whether one could cash out such conditions without adverting to (temporal) representational content. This might not in itself be objectionable. Recall that Burge denies one can characterize in non-representational terms what perceptual constancies – the marks of perception – are. But it would limit the informativeness with which one could answer the 'origins' question.

On the other hand, if (v) is rejected, then we do not have even that much of an answer. Perhaps we might approach the matter by identifying – but then

removing – the conditions that render the ascription of temporal content explanatorily otiose in our simple cases. With our organisms, the time-course from perception to use is fixed and so short relative to the time-scale of ethologically relevant environmental variations as to be negligible. Further, there is but one representation at a time poised for use. Suppose, however, that this were not so: perceptual representations, or states storing their contents, could be used at various later times, and multiple such states might be poised for use. This could benefit the organism if differences in the time of use and differences in which representation was used now could affect ethological success. Scrub jays provide a particularly spectacular example. They remember, not only the locations of tens of thousands of caches, and not only what kind of food has been cached, but also *when* each was cached – as is demonstrated by their preference for caches whose contents haven't decayed to the point of inedibility (which requires them to learn as well the differing rates of decay for different kinds of food). (Clayton and Dickinson, 1999) Carrying the contents of perceptions forward in time to be used differentially according in part to when the perceptions occurred plausibly suffices for the attribution of temporal content to the representations themselves.

But to identify a capacity sufficient for temporal representation in perception is not necessarily to identify its origins. Something more minimal might likewise suffice. Earlier, we noted Burge's suggestion that the perception of change or movement is the simplest sort of temporal representation in perception. Here too we have a feature not found in our simple organisms: if we no longer restrict ourselves to organisms capable of perceiving only lightness, the possibility of explanatorily useful ascriptions of temporal representation in perception again can arise.

So which candidate for the origin of temporal representation in perception should we favor: the capacity to carry forward the contents of perceptions for use at various times or the capacity to perceptually track change or movement? Or perhaps there may be some other deployment of temporal representation in perception even more minimal in its conditions than these. I doubt that reflection on the nature of perception and temporal representation can enable us to answer such questions. Further illumination, especially concerning the phylo- and ontogenetic questions, must emerge from the 'rough and tumble of explanation' (p. 57) as pursued in the empirical sciences.¹⁵

Department of Philosophy
Johns Hopkins University

NOTES

¹ One methodological remark. A capacity for temporal representation is arguably ubiquitous among *actual* organisms capable of perception (Burge, 2010, p. 526 – henceforth, unadorned page references are to this text). So, perhaps it's no surprise that possible, but non-actual creatures might arise in a challenge to Burge's *constitutive* claim – where constitutive claims, for Burge, concern conditions on something's being what it is and ground explanations of its nature (p. 5). That said, the creatures that come up in our discussion seem *prima facie* *nomologically* possible. If they are not, perhaps those uncomfortable with broader modalities and 'natures' might allow that sometimes indulging in speculative science fiction and wallowing in marginal cases can illuminate a subject, even if only by forcing us to articulate – and perhaps disentangle – otherwise inchoate strands of our fundamental conceptions. In any event, Burge is not skeptical, and the possible creatures discussed are either his own or minor variants thereof.

² Burge's reasons are not directly germane to our topic, but the basic idea is that attribution in perception always guides singular application, whereas in conception it needn't. In his view, attribution that does not guide singular application marks a kind of 'freedom from the here and now' possible in propositional thought and inference. It is reflected in propositional structure by the existence of predication, the kind of attribution distinctive of conception (pp. 539–43).

³ Because Burge maintains that correctness conditions are constitutively normative, it follows, given the distinctions drawn in the last two paragraphs, that there are non-biological, natural norms: norms in nature that can obtain and apply without awareness or understanding on the part of any subject but that cannot be reduced to the biological (in particular, can't be cashed out in evolutionary terms).

⁴ This, if correct, undermines Schulte's (2015) attempt to defend teleosemantics from Burge's 'breadth of application' objection by incorporating Burge's own considerations concerning perceptual constancies. Schulte assumes that an intended reduction of intentionality can appeal to the notion of perceptual constancy without circularity. He does not mention Burge's claim to the contrary.

⁵ However, even if the naturalistic acceptability of intentional notions does not require relating the intentional to the non-intentional, it may be that it *is* required if we are to *unify* explanations that deploy intentional notions with those that do not. Cf. Gross, 2005.

⁶ One might even wonder what that would amount to. But care must be taken that any charge here of unintelligibility doesn't turn on a questionable folk constraint on perceptual content. We might not naturally conceive of ourselves as causally interacting with times or temporal features in a way that enables us to perceive them. But, on Burge's view, what matters is rather whether there are good grounds, given what explanation requires, to maintain that, in perception, temporal features are attributed as a result of the relevant kind of computational process. Cf. Burge's critical discussion (p. 444) of Spelke's (1988) appeal to the 'abstract' in arguing against the representation of objecthood in perception. Moreover, the perception of temporal features can involve integrating information from multiple modalities that transduce signals at different speeds, a process that arguably sustains temporal constancies. That said, it's a common view that temporal representation, including in perception, crucially differs from perception of other features. Premise (v) in effect sums up Burge's attempt (pp. 521–6) to properly place the origins of this distinctive representational category in his overall framework.

⁷ Note that temporal sensitivities are not, or at least need not be, sensory registrations of information, since they needn't involve the senses: circadian rhythms, for example, need not be set or calibrated via sensory registration (though they often are). But temporal sensitivities are *analogous* to sensory registrations of information in not sufficing for representation. They are analogous also in requiring a relation to mechanisms sustaining perceptual constancy in order to rise to the level of representational content. They are disanalogous, however, in that the relation may be indirect.

⁸ Burge's discussions of *specific kinds* of perception – an animal perceptually tracking a moving particular, a bee associating a phase in a circadian cycle with a perception of sugar water so as to return for a drink the same time the next day (pp. 521–2) – provide examples of less minimal, and perhaps more obvious, ways that perception relies on temporal sensitivity. One might wonder whether, for the general case, Burge could appeal to the harnessing of an even more minimal temporal sensitivity – for example, just the correlation of perception with the present – and thus avoid bringing memory into his argument at all. (Cf. also Burge (p. 522) on the correspondence of temporal sequences of sensory registrations of information and temporal sequences of their causes.) Burge might reply that this would not suffice because a capacity for temporal representation requires a capacity for representing a temporal *framework* (cf. the language on p. 526, and Burge, 2011). Perhaps something parallel is required also for a temporal sensitivity. Incidentally, because Burge's argument does not turn only on this simpler correlation, the kind of harnessing on which it relies arguably also involves a second sort of indirectness beyond the harnessing to perception of something else: it adverts to the way temporal sensitivity is implicated in the mechanisms of memory required for the use of perception, as opposed to temporal sensitivity being implicated directly in the mechanisms sustaining a constancy. But that the argument is so constructed does not preclude more direct harnessing of this second kind.

⁹ Burge's simple organisms show that – unlike Ganson *et al.* (2014) and Dretske (1988) – he doesn't require flexible behavior in order for there to be representational content. Ganson *et al.* argue that, where there is inflexibility, the ascription of content is explanatorily otiose. Burge can reply by invoking its role in mechanisms sustaining perceptual constancies. (To be sure, Ganson *et al.*, earlier in their paper, raise objections to Burge's views on perceptual constancies and their relation to perception. But their remarks about flexibility are presented as a further, independent objection.) Incidentally, Burge's simple organisms would lack as well a capacity to perceive objects as such.

¹⁰ One might similarly question (i)'s requirement that the use be *representational*. Maybe *any* use would satisfy the *rationale* for (i), regardless of whether it was representational. It's unclear, however, what a non-representational use would be if – as Burge suggests – use of a perception *engenders* representationality in the use.

¹¹ It might be suggested that the perception of change is only the 'simplest sort of temporal representation *in perception*' (italics added), not necessarily the simplest sort of temporal representation, so that Burge's remark leaves room for simpler organisms capable of perception and thus of temporal representation, but not of temporal representation *in perception* – for example, creatures without temporal representation in perception that nonetheless meet the constitutive condition for perception by being capable of temporal representation in perceptually-based memory or elsewhere. But, even if this were right, it would remain the case that the examples Burge discusses in support of premise (v) seem not to cover cases simpler than those involving perception of change. In any event, it's not clear that this *is* right, at least by Burge's lights, since he holds (cf. fn. 13 below) that all perceptual content is temporally indexed, so that there can't be an organism capable of perception that lacks temporal representation *in perception*. (Arguably, there can't *be* perceptually-based temporal representation in memory without such temporal indexing in perception. The functional relations that enable a memorial representation as of, for example, [that, red, *then*] might require that there have been perceptual representations as of [that, red, *now*].) This raises the question: why isn't this temporal indexing, by Burge's lights, an even *simpler* sort of temporal representation *in perception* than perception of change? Perhaps Burge might reply that any perception involves a capacity to perceive change even though it need not involve a capacity to perceive change as such – and that, in his remarks on simplicity, he intended only the former. This would not map naturally onto the language he uses on pp. 528–9.

¹² It is unclear whether on Burge's own view, according to which our bead on both the form and content of representations is given by their explanatory work, content *per se* is a limited resource.

¹³ Though he doesn't *argue* for it in this manner, Burge does *endorse* the claim that all perceptual tokens are egocentrically indexed to the present. He thus maintains, not only that a *capacity* for temporal representation is constitutively necessary for a *capacity* for perception, but also that each *token* perception necessarily involves temporal representation. This token claim, though mentioned in *Origins* (p. 525), is not prominent there. But see his Dewey Lectures (Burge, 2011). The claim may seem in tension with his remarks, quoted above, that '[a]n animal could have a visual system capable only of attributing lightnesses or colors' (p. 496) and '[p]erhaps there could be a perceiver capable only of awakening and having a perception of color or shape' (p. 528). But the first quote refers only to the attributive aspect of the organism's perceptions, whereas the temporal index is *de re* (p. 521). The second quote might likewise be understood. (That Burge does *not* argue merely by appeal to perception's functioning to register the 'here and now' might also reflect the need to establish a capacity for representing a temporal *framework*, as per my fn. 8 above.)

¹⁴ One strategy worth mentioning – because it draws on Burgean material – would deploy aspects of his anti-individualism. Perhaps representing temporality might be necessary for having any *veridical* perceptions, which in turn might be necessary for having states with representational content at all. For example, slightly more spelled out, perhaps:

- (1) Perceivable objects do not a-temporally possess properties.
- (2) If perceptual states did *not* temporally represent, they would therefore be non-veridical.
- (3) For an individual to have *any* representational state, the state must be associated with some appropriate *veridical* representational state. (Burge's principle (B), p. 68.)
- (4) Some of those states must be perceptual (if there is to be sufficient determinacy of content).
- (5) So, if perceptual states did not temporally represent, we would not have *any* representational states.
- (6) But we do have some representational states.
- (7) So, we have some perceptual states with temporal content.

This particular argument would have to overcome a variety of *prima facie* obstacles. For example, first, what precludes a type of attribution that is neither temporal nor a-temporal, or is non-specific between them – or that requires temporal features in its circumstances of evaluation without *representing* them (as some might argue is the case with non-modal claims and actuality)? Second, Burge's principle (B) does not require that it be the individual herself that be in the appropriate veridical representational state: it could be an ancestor who lived in a radically different environment; so, how may we conclude (as seems to be the intent of Burge's own argument) that a capacity for perception constitutively requires a capacity for temporal representation in that very individual? Third, why should Burge's principle (B) – properly construed – require an associated representational state that is veridical in *all* its aspects? A perception as of a certain-sized red rectangle might be veridical with respect to color – thus satisfying the conditions necessary for sufficiently determinate color content – while yet misrepresenting size and shape.

¹⁵ A version of this material was presented at the Burge Workshop sponsored by the Norwegian University Center in St Petersburg, Russia. Many thanks to the organizer Carsten Hansen. For helpful discussion, thanks to Ned Block, Sue Carey, Georges Rey, and, especially, Tyler Burge. Thanks also to Justin Bledin and an anonymous referee for very helpful written suggestions.

REFERENCES

- Brandom, R. (1994). *Making It Explicit*. Cambridge, MA: Harvard University Press.
- Burge, T. (1991). "Vision and Intentional Content," in E. Lepore and R. Van Gulick (eds) *John Searle and His Critics*. Oxford: Blackwell, pp. 195–213.

- Burge, T. (2003). "Psychology and the Environment: Reply to Chomsky," in M. Hahn and B. Ramberg (eds) *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge, MA: MIT Press, pp. 451–470.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Burge, T. (2011). "Self and Self-Understanding," *The Journal of Philosophy* 108, pp. 287–383.
- Chomsky, N. (1995). "Language and Nature," *Mind* 104, pp. 1–61.
- Chomsky, N. (2003). "Internalist Explorations," in M. Hahn and B. Ramberg (eds) *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge, MA: MIT Press, pp. 259–288.
- Clayton, N. S. and Dickinson, A. (1999). "Memory for the Content of Caches by Scrub Jays (*Aphelocoma coerulescens*)," *J Exp Psychol Anim Behav Process* 29, pp. 82–91.
- Davidson, D. (2001). *Subjective, Intersubjective, Objective*. Oxford: Oxford University Press.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- Fodor, J. (1990). *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Gallistel, R. (1990). *The Organization of Learning*. Cambridge, MA: MIT Press.
- Ganson, T., Bronner, B. and Kerr, A. (2014). "Burge's Defense of Perceptual Content," *Philosophy and Phenomenological Research* 87, pp. 556–573.
- Gross, S. (2005). "The Nature of Semantics: On Jackendoff's Arguments," *The Linguistic Review* 22, pp. 249–270.
- Millikan, R. G. (1989). "Biosemantics," *The Journal of Philosophy* 86, pp. 281–297.
- Schulte, P. (2015). "Perceptual Representations: A Teleosemantic Answer to the Breadth-of-Application Problem," *Biology and Philosophy* 30, pp. 119–136.
- Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Spelke, E. (1988). "Where Perceiving Ends and Thinking Begins: The Apprehension of Objects in Infancy," in A. Yonas (ed.) *Perceptual Development in Infancy*. Hillsdale, NJ: Lawrence Erlbaum, pp. 197–234.
- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.