Noname manuscript No.
(will be inserted by the editor)

Strong Liberal Representationalism

(Pre-print. Forthcoming in *Phenomenology and the Cognitive Sciences)*

Marc Artiga
Universitat de València

Abstract The received view holds that there is a significant divide between full-blown representational states and so called 'detectors', which are mechanisms set off by specific stimuli that trigger a particular effect. The main goal of this paper is to defend the idea that many detectors are genuine representations, a view that I call 'Strong Liberal Representationalism'. More precisely, I argue that ascribing semantic properties to them contributes to an explanation of behavior, guides research in useful ways and can accommodate misrepresentation.

1 Introduction

Which states should qualify as genuinely representational? How many organisms can be truly said to possess states with semantic properties? Consider the following examples: some anaerobic bacteria possess organelles called 'magnetosomes' that orient themselves along the Earth's magnetic field and aid bacteria to move to deep and relatively oxygen-free sea levels; the Venus flytrap (*Dionaea muscipula*) catches its prey (mostly insects and spiders) by closing its mouthlike leaves when some tiny hairs ('triggering hairs') on their inner surfaces are touched upon; some mosquitoes (e.g *Aedes aegypti*) that have evolved a preference for human hosts find their victims by being attracted by some compounds present in human sweat, such as lactic acid, ammonia, ketones or sulfides as well as by CO_2 and heat. Are these behaviors driven by internal representations?

The main goal of this paper is to argue that these examples as well as many other internal states of animals, plants and bacteria are genuine representations. More precisely, I will focus here on a set of states that have been labeled 'receptors', 'C-states' or 'detectors'. *Faute de mieux*, here I will employ

Address(es) of author(s) should be given

the latter.[1] Detectors can be defined functionally: on the input side, they are stimulus-specific, in the sense that they are triggered by a particular kind of input. For instance, exteroceptors (such as thermal nociceptors, photoreceptors or Merkel cells) are detectors set off by a specific stimulus from the external world (extreme temperature, light and pressure, respectively). Other detectors are used in proproception or interoception to sense the internal states of the organism, such as muscle length or joint angle (see Kandel et al. **?**, p. 475-80; Ganson, **?**). In some cases, detectors might be triggered by other detectors, rather than by external or internal cues. On the output side, detectors produce very specific effects. As some of the examples will show, often this effect is some form of behavior such as flying away, producing certain sounds or approaching an object. In other cases, the particular effect merely consists in triggering a specific pattern of neuronal activity downstream. Finally, detectors tend to be strongly phylogenetically determined, i.e. they are largely innate. This last property is less central than the other two, though it might contribute to describing the kind of example that disputants tend to have in mind (and will be relevant in some of the arguments to come, especially in section **??**). To make the notion of detector clearer, let me illustrate it with some examples.

Nocturnal crickets are often preyed by bats, which use echolocation to locate their victims. As a result, the cricket *Gryllus bimaculatus* (among other insects) has evolved auditory organs with a bat-detecting function. The cricket's eardrums are located on the tibiae of the fore legs. The first neural relay is constituted by 60-70 auditory receptors; about one-fourth of them preferentially fire in response to high frequencies (including ultrasounds) and the remaining auditory receptors respond to low frequencies, which are probably used in intraspecific communication. When the auditory neurons attuned to ultrasounds fire, this information is processed primarily by two kinds of interneurons that receive monaural excitatory inputs from receptors. One of them is called 'Ascending Neuron 2' (or AN2). According to **?**, p. 1, "AN2 firing is necessary and sufficient to trigger avoidance steering". Thus, there is a reliable causal path: $bat \rightarrow ultrasond \rightarrow eardrum \rightarrow auditory\ receptor \rightarrow AN2 \rightarrow steering\ away$. Auditory receptors and ascending neurons are clear examples of detectors in our sense: they are triggered by a very specific stimulus, generate a distinct effect and are phylogenetically determined. Other paradigmatic detectors with similar features are easy to find in many species: in the silkmoth *Bombyx mori*, for instance, the presence of a single pheromone component, bombykol, elicits the full sexual behavior in males (**?**). Likewise, fire ants (*Solenopsis invicta*) carry the corpses of other members of the colony from the nest when they detect certain chemical cues (Choe et al.2009, **?**). The examples mentioned in

[1] 'Receptor' is used for certain biological kinds that might or might not coincide with the set of states that are supposed to be captured by the philosophical use of the term, so the use of this expression is a potential source of misunderstandings. Furthermore, this term misleadingly suggests that these states are purely passive but, as I will argue below, they are also defined by the specific behavior they trigger. 'Detector' is not ideal either, since this expression seems to beg the question in favor of representationalism. Finally, 'C-state', which has also occasionally been employed in the literature, is not the most elegant proposal and is much less evocative than the other two.

the first paragraph involving bacteria,[2] Venus flytraps and mosquitoes clearly qualify as well. Humans also probably posses detectors, such as sensory receptors. Thus, the study of detectors spans a wide variety of scientific disciplines such as microbiology, plant neurobiology, cognitive ethology or neuroscience.

Now, are detectors representations? An affirmative answer leads to Liberal Representationalism (LR):

Liberal Representationalism (LR) Many detectors[3] are representations.

Why is LR formulated in terms of 'many detectors', rather than all of them? One reason is that this is probably the strongest version of a liberal view that can sensibly be defended. After all, some detectors might play a very special role in some very particular systems that cannot be described as truly representational.[4] Nature is complex and messy and universal claims rarely hold. In any case, LR should be interpreted as the claim that *typically* or *generally* detectors are representations.

Now, most philosophers who have addressed this question reject LR (see, for instance, ?; ?; ?; ?; ?;?; ?; ?; ?; ?; ?; ?; ?; ?; ??; Cf. [author]). Note that detectors look very different from those states that have traditionally been regarded as paradigmatically representational, namely beliefs and desires. The belief expressed by the sentence 'It rained yesterday' can be triggered by a variety of situations (I can see some puddles, hear the rain or hear someone saying 'It rained yesterday', for instance). Likewise, it can have many different effects, from causing further beliefs (e.g. 'It has been a rainy week', 'The pavement is probably wet',...) to producing a variety of behaviors (e.g. to take an umbrella, to stay at home,...). Finally, beliefs and desires are typically not innate.

Even among those who accept LR (for reasons that will be clear in the next section) it is common to regard detectors as a limiting case of representation "in the way zero is a number" (Millikan, 2009: 406). In other words, although they might accept that detectors count as representations, this is an unintended consequence of a theory that has been developed to explain more traditional kinds of representational states. It is assumed that classifying detectors as representations is a theoretical cost for the theory, a bullet they have to bite, so they tend to stress that detectors lack some important features possessed by typical representations, which show them to be marginal cases of intentionality (Dretske, ?; Millikan, ?).

[2] For a discussion of the adaptationist explanation of bacteria magnetotaxis, see O'Malley, ?: 29-38.

[3] Strictly speaking, the claim is that certain activity patterns of detectors are representations. For simplicity, however, I will often claim that detectors are representations. When the distinction between mechanisms and activities is relevant, I will explicitly distinguish the two.

[4] Let me stress that I do not have an particular example in mind. In any case, all examples I will discuss here qualify as truly representational. I would like to thank a reviewer for pressing me on this point.

In contrast to this perspective, in what follows I would like to argue that detectors are full-blown representational states. This is the view I label 'Strong Liberal Representationalism' (or SLR):

STRONG LIBERAL REPRESENTATIONALISM (SLR) Many detectors are genuine representations.

Where by 'genuine' I mean not marginal or a limiting case, but a full-blown instance of representation. In other words, SLR emphasizes the independent plausibility of the idea that detectors are representations, by showing that they meet the main conditions for a state to qualify as a standard case of representation. As a result, accepting detectors within the category of representations is not a bullet to bite or a theoretical cost, but a plausible claim that any approach should try to accommodate. Furthermore, I will argue this claim can be defended independently of any commitment to a particular naturalistic theory. As a result, if the arguments in favor of SLR developed in this paper are sound, the default assumption in assessing any theory should be that detectors are indeed representational states.

1.1 SLR: First pass

If one seeks to address the question of whether a certain entity is a representation, here is a sensible strategy: consider our best theories of what representations are and assess whether this entity fulfills the conditions for qualifying as such. There are certainly a number of theories on the market, but in I think many of the classical approaches will deliver the same verdict: detectors do meet the conditions for qualifying as representations.

Consider, for instance, one of the most popular approaches: Teleosemantics. Roughly, according to a classical version of teleosemantics, a representation is a state S produced by some mechanism (M_1) and used by an other mechanism (M_2), in which both mechanisms have certain functions. On the one hand, the sender's function is to emit a state R when some event S holds. On the other, the mechanism using R has the function of producing certain response E, which in normal conditions is only successful because S (the referent) is the case[5] (Millikan, ?; Price, ?). Finally, some authors add that R (or the set of representations) carries correlational information about S (or a set of states) (Shea, ?; Martínez, ?).

[5] Millikan (?) emphasizes the idea of a systematic mapping relationship between a set of representations and a set of entities. So a more precise description of her view would say that the function of the sender is to emit a range of states which are supposed to map onto a range of states according to a certain mapping function, and which might elicit a range of behaviors. Nonetheless, this criterion can be satisfied by a mechanism that produces only two states and behaviors: a cell firing (or not firing) at a time t and place p representing the presence (or absence) of a predator at t near p, which triggers an evasive behavior (or it does not). Consequently, these apparently more stringent conditions can easily be met by detectors as well.

The key point we are mostly interested in is one on which there is large agreement: this classical teleosemantic approach entails that many detectors are representations. Consider one of our previous examples: Crickets' AN2 interneurons (M_1) have the function of firing (R) when there is a bat around (S) and in normal conditions what explains the successful performance of the motor system's (M_2) function (E) is the presence of a bat. Furthermore, activity in AN2 interneurons carries correlational information about bats, so all conditions of teleosemantics are fulfilled. This result is likely to generalize to most detectors: detectors are mechanisms that have been selected for generating an internal state whenever certain event takes place, and they are supposed to produce a specific effect, whose success depends on this event. Thus, teleosemantics seems to entail STRONG LIBERAL REPRESENTATIONALISM.

Unsurprisingly, there is wide agreement on the idea that naturalistic theories tend to have this consequence (Burge, **?**; Butlin, **?**; Ramsey, **?**; Schulte, **?**). However, most philosophers view this entailment as a bug rather than a feature, and adopt different attitudes as a response to it. Some people take this result as a *reductio* of teleosemantics and other naturalistic theories (e.g. **?**), whereas others modify their preferred naturalistic theory in order to avoid SLR (Butlin, **?**; Neander, **?**; Schulte, **?**). In this case, the challenge is to find some feature related to the nature of representations that distinguishes detectors from full-blown representations.

In my opinion, the fact that SLR follows from some of our best naturalistic theories of content should be taken as evidence for SLR. Nonetheless, I also think more needs to be done to make a convincing case for SLR. For one thing, the conclusion of an argument can only be as strong as the weakest of its premises, so although this argument might put some pressure on teleosemanticists and other naturalists, those who reject these frameworks will not be moved by these considerations. Moreover, some naturalistic theories have been developed to block this entailment, so we need to find some positive reasons for favoring a liberal approach. Finally, the fact that SLR follows from some of our best naturalistic approaches does not say anything about the plausibility of SLR as such. To address these concerns, in this paper I would like to discuss whether SLR is a reasonable claim, irrespective of the support it can gather from other proposals, so I will rely as little as possible on the truth of teleosemantics or any other naturalistic approach. In particular, I am interested in assessing whether detectors play the explanatory roles usually associated with representations.[6]

In the following sections I will defend STRONG LIBERAL REPRESENTATIONALISM by paying attention to some central properties of representations.

[6] The fact that scientists often *assert* that detectors are representations could be taken as an additional piece of evidence in favor of SLR. Figdor (**?**), for example, has interestingly argued that we should interpret these assertions literally. Although these results are sympathetic to the main conclusion of this paper, the arguments differ: here I would like to focus on the explanatory role of SLR, rather than on the linguistic interpretation of scientific claims.

More precisely, I will concentrate on three features that are usually regarded as central and paradigmatic of representational states:

C1 An appeal to representations contributes to an explanation of behavior.
C2 Attributing representational content guides research in useful ways.
C3 Representations can be false (they can misrepresent).

The rest of the paper is devoted to showing that detectors fit this characterization. In section **??** I will defend that ascribing semantic properties to detectors contributes to explaining behavior, by discussing three different models. This is the claim I will defend in more detail, partly because there are different ways of understanding how representational content explains behavior and partly because most critics of SLR have focused on this feature. In section **??** I will discuss how these representational attributions also guide research in useful ways. In section **??** I will argue that detectors can misrepresent and in the final section I will address some objections.

2 Explaining Behavior

A central aspect of our concept of representation is that an attribution of semantic content contributes to an explanation of behavior (**?**, p. 280; **?**, p. 22; **?**, p. 498). We explain that Carl von Linné went to Lapland by appealing to his belief that there were interesting plants, animals and minerals in this region and his desire to collect and study them. As Ramsey (**?**, p. 129) suggests:

> A minimal requirement for a successful functional specification of any notion of representation is that the content – or, if you like, the fact that the representation *has* a semantic content– be an explanatorily relevant fact about that state.

Most philosophers who have addressed this question think that assuming that detectors have semantic content does not contribute to any satisfactory explanation. Thus, if the explanatory role of the detector's content can be vindicated, it will undermine one of the main objections against Strong Liberal Representationalism and will provide an important argument in favor of the claim that detectors are genuine representations.

Unfortunately, there is a difficulty lurking ahead: even though most philosophers agree that an appeal to semantic properties is explanatory, they strongly disagree about *why* this is so. More generally, the debate on the nature of representations and their explanatory role has taken place within the framework of a realist conception, according to which mental representations are physical particulars that interact causally in virtue of non-semantic properties, but in ways that are faithful to their semantic properties. Thus, a very controversial philosophical question has been how to account for the explanatory role of content, given that an account in terms of physical particulars within the system seems to provide a complete explanation of behavior.

How can a liberal representationalist address the alleged explanatory idleness of content attributions to detectors, given this lack of consensus? The strategy I will pursue is to discuss three different models of why semantic content explains behavior and argue that according to a reasonable interpretation of all of them the content of detectors is explanatory. In particular, I will focus on Dretske's approach, for being one of the most influential accounts, Shea's model for providing a recent, clear and plausible theory focused on sub-personal states, and finally on Ramsey's proposal, for his being a staunch critic of LIBERAL REPRESENTATIONALISM.

2.1 Dretske's approach

How do semantic properties contribute to explaining behavior? **?** provides one of the most prominent answers. To understand his model, however, we need some stage setting. First of all, according to Dretske a behavior is a process: the causing of a body movement (E) by some internal state (R) within a system. Thus, the *explanandum* of a representational attribution is not, say, why my hand moved (E), but why an internal state of my brain (R) caused my hand to move (E). Secondly, Dretske draws a distinction between two kinds of causes, that he labels 'triggering' and 'structuring' causes. A triggering cause is an event that causally explains why a particular event occurs at a certain point in time. In the context of behavior, a triggering cause of my behavior is an event (e.g. the presence of a fly near my face at a particular time) which is causally responsible for the occurrence of the internal state (R), which in turn causes a hand movement (E). In contrast, a structuring cause is an event that is causally responsible for the connection between two events in a system, such as R and E. In other words, a structuring cause explains why the mechanism is wired in such a way that an activation of R elicits E. Thus, whereas a triggering cause explains why R occurs *now*, a structuring cause can explain why *R causes E*.

Now, on Dretske's view semantic properties are structuring causes of behavior: the fact that a given representational state means S accounts for the fact that an internal state R causes a movement E. Imagine that R is some internal state of an organism, which is usually produced when an event S holds. As a result, R indicates or carries correlational information about S. Suppose further that there is some movement E that would be beneficial for the organism to perform when S holds. If, through some recruiting process (e.g. learning, natural selection) internal state R becomes linked to movement E because often enough R correlates with S, then the R becomes a representation and the fact that it indicates S becomes a structuring cause of R *causing* E. According to **?**, this is a mechanism by means of which content can become a structuring cause of behavior.

To illustrate the idea, let us try to apply this model to crickets. When a bat is nearby (S), a cricket obviously benefits from flying away (E). As a result, through natural selection an internal state (i.e. activity in ascending neurons

AN2) became highly correlated to the presence of the bat and also wired to an evasive movement. Thus, activity in AN2 interneurons tracks bats being around and the fact that this internal state has that content is a structuring cause of interneurons causing the cricket's flight. Thus, if Dretske is right about the explanatory role of semantic properties, the representational content of AN2 interneurons is a structuring cause of behavior.[7] It is not hard to see that the same reasoning can be extended to other detectors.

Interestingly, Dretske denies that the semantic properties of detectors are explanatory (**?**, p. 93).[8] More precisely, he holds that only the semantic properties of representations recruited by some learning process explain behavior. This would exclude many detectors, since, as we saw above, they typically are phylogenetically determined, rather than learned. Yet I think his arguments for a clear distinction between the explanatory role of learned and unlearned representations can be questioned on different grounds.

First of all, he follows **?** and **?**, p. 147-152 in distinguishing selectional from developmental explanations and argues that selectional explanations cannot explain why organisms have certain features. Evolution can account for the fact that a particular trait has certain distribution within a population, but not how it originated:

> As Cummins (1975) notes, natural selection (assuming this is the chief pressure for evolutionary change) does not explain why organisms have the properties for which they are selected any more than Clyde's preference for redheads explains why Doris, his current favorite, has red hair.[9](**?**, p. 92)

Now, if natural selection cannot explain the existence of any trait, then it cannot vindicate the idea that R indicating S in the evolutionary past explains the fact that in current organisms R causes a movement E. If natural selection cannot explain why organisms have certain trait, then *a fortiori* it cannot explain behavior, which is of course a trait of organisms. (**?**, p. 92-95).[10]

Now, there are two main reasons why this reasoning fails to jeopardize the explanatory role of unlearned representations. First, the main premise of this argument has been largely discredited: nowadays most philosophers are

[7] Let me hasten to add that questions about the indeterminacy of content are irrelevant here. Certainly, the behavior could be described in various ways (e.g. flying away from bats, avoiding a predator, evading a deadly flying mammal, etc..) and, accordingly, there might be some content indeterminacy (e.g. bat, predator, deadly flying mammal, etc..). However, in any of these descriptions, the ascribed semantic content is equally explanatory.

[8] Nonetheless, he accepts that they are still representations. As a result, he might be taken to embrace LIBERAL REPRESENTATIONALISM, but reject STRONG LIBERAL REPRESENTATIONALISM.

[9] I apologize for reproducing a sexist example.

[10] Note that this idea is in tension with Dretkse's (1988, p.91-92) suggestion that indication in the evolutionary past can qualify as a structuring cause of behavior. If an entity M is a structuring cause of R *causing* E, then presumably M explains (or at least contributes to an explanation of) R causing E. As a consequence, M explains a property of organism. Thus, the role of past indication as a structuring cause is hard to reconcile with the idea that natural selection can not explain why organisms possess certain features.

happy to accept that natural selection can indeed explain why certain kinds of traits originate, as well as their distribution within a certain population (**??**).[11] Natural selection can account for the existence and design of eyes, not just for their universal distribution among humans.

Secondly, Dretske's model does not actually require that the recruiting process explain the existence of a trait: it suffices if it accounts for its maintenance (see, Dretske, 1988, ft. 10, p. 89) and natural selection can surely account for the stabilization of a certain trait within a population of organisms. As a result, even if the content of innate representations has been fixed phylogenetically, they can explain why current organisms are wired in such a way that R causes E.

Dretske provides another consideration for resisting the idea that the content of innate representations can explain behavior: he claims that the behavior caused by an unlearned representation R does not use R's informational content because it is not sensitive to it. More precisely, modifying the informational content carried by R fails to produce any change in behavior:

> Even if through a recent freak of nature (recent enough so that selectional pressures had no time to operate) the occurrence of [R] in contemporary moths were to signal not the approach of a hungry bat but the arrival of a receptive mate, [R] would still produce [E] - would still produce the same evasive flight maneuvers (Dretske, 1988, p. 93)

The proviso, of course, is crucial: '*recent enough so that selectional pressure had no time to operate.*' I think the need for this proviso shows why the argument fails. Let me explain.

According to Dretske, the fact that innate representations indicated such and such does not explain behavior. At the same time, however, he admits that R's past indication of S is a structuring cause of behavior (i.e. of R causing E), so if past instances of R had not indicated S, then R would not currently cause E. How can Dretske maintain that indication is a structuring cause of behavior and, at the same time, that changing facts about indication would not modify behavior? By considering a change that is so recent that no suitable recruiting process could have taken place. If we consider a behavior that has been recruited by natural selection and consider a change in the indication profile that does not give enough time for natural selection to operate, then this change in indication does not cause a change in behavior. However, the same is true of any recruiting process – learning included (see Dennett, **?**, p. 122). Imagine, for instance, that moths *learned* to avoid bats by relying on AN2 neurons, instead of this behavior being innate. The same quote would be true if 'selection pressure' is substituted by 'learning':

[11] In a nutshell, here is the reasoning: "natural selection can reshape a population in a way that makes a given variant *more likely to be produced* by the immediate sources of variation than it otherwise would be. As selection changes the background in which mutation and recombination operate, it changes what those factors can produce" (**?**, p. 29). The idea that still remains controversial is whether natural selection can explain why a particular individual possesses a certain trait.

> Even if through a recent freak of nature (recent enough so that *learning* had no time to operate) the occurrence of [R] in contemporary moths were to signal not the approach of a hungry bat but the arrival of a receptive mate, [R] would still produce [E] - would still produce the same evasive flight maneuvers.

More generally, the fact that this counterfactual holds does not show anything about the explanatory role of indication, because it depicts a case in which the recruiting process that is necessary for the cause to change the effect is missing. Any cause could be found to be irrelevant for its effect if one considers a situation in which the cause is modified but not enough time is given for the channeling conditions to operate.

As a final consideration against Dretske's view on the relevance of the distinction between learning and natural selection, the idea that the semantic properties of innate representations fail to explain behavior goes against most research in cognitive science. Whether a representation is innate or learned is not usually taken as a piece of decisive evidence for establishing whether its semantic properties of explanatorily relevant or not (see, for instance, Cummins, **?**; Horgan, **?**).

Summing up, I think Dretske fails to provide convincing reasons for resisting the idea that the semantic properties of detectors explain why internal states are linked to certain behavioral responses. Therefore, if Dretske's model of behavioral explanation is adopted, then the detector's semantic properties are explanatory. Of course, Dretske's approach has been criticized on various grounds and it would be hazardous to base the vindication of STRONG LIBERAL REPRESENTATIONALSIM on a single model. To circumvent these worries, let me discuss a plausible and more recent approach that focuses on subpersonal states: **?**'s model.

2.2 Shea's approach

? provides a powerful and empirically oriented account of how semantic properties are explanatorily relevant that differs from Dretske in some respects. On the one hand, Dretske assumed that semantic properties explain behavior, understood as the process that consists in an internal state causing a bodily movement. Shea, in contrast, holds that representations primarily explain *successful* behavior, which is understood as a world-involving effect of a system (**???**). His model tries to understand, for instance, why holding a true cognitive representation about a ball approaching my body explains why I successfully caught it. In Godfrey-Smith's (**?**) terms, true representations are a 'fuel for success'.

Shea's model is easier to understand with an example that he provides (**?**, p. 82-84). Suppose we observe a robot moving along a line towards a certain point T, slowing down as it approaches it and stopping when it reaches T (see figure **??**). Imagine we want to explain how this machine managed to successfully reach T and so we open the box and investigate how its insides work.

Suppose we see that activity in the internal component r reliably correlates with the distance of the machine from the origin. A second component δ registers the distance from T, and this correlation is achieved by subtracting the activity of r form a fixed level of activity t. On the other hand, component a registers the velocity of the wheels and causes the robot to move. A monotonic transformation from δ to a controls the machine's velocity and drives it to point T. Finally, suppose that some stabilizing process has endowed the system's components with functions (e.g. suppose there is a power source at T and the fact that the machine reached point T in the past has contributed to the persistence of the system -?, p. 62). According to many, that would be enough r to stand for the distance from the origin, for δ to represent the distance to T and for a to be about the velocity of the wheels.

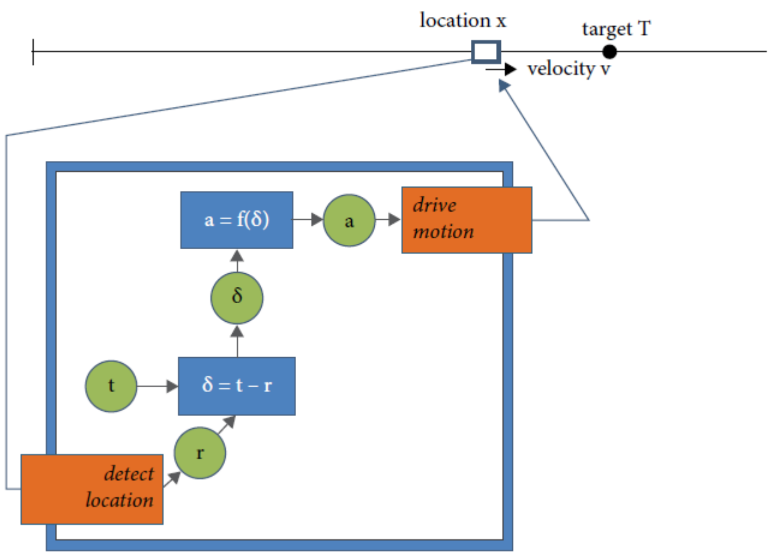

Fig. 1 This system is disscused in the main text. From **?**, p.67

On **?**'s approach, the semantic properties of any of the robot's internal components explain successful behavior because the system has evolved to exploit the correlational information that they carry about different variables (distance from origin, velocity of the wheels, etc.) in order to successfully reach point T. A satisfactory explanation of how the robot reaches point T has to appeal to the fact that each of these components tracks a different feature of the world. Internal processing over these elements constitutes an algorithm for successful performance of this world-involving task (**?**, p. 82).

Again, let us apply this model to crickets. Scientists wondered how crickets manage to successfully escape from bats and when they opened the 'cricket box', they found at least two components which turned out to be essential. On the one hand, the first neural relay is constituted by auditory receptors that carry correlational information about ultrasounds. The second component is AN2 interneurons, which correlate with the presence of bats and whose activity triggers an evasive response. Thus, it seems that the system is exploiting the information that auditory receptors carry about ultrasound waves and the information that AN2 neurons carry about bats in order to successfully steer away from them. As before, the explanation of how the system (crickets) successfully performs a world-involving behavior (avoiding bats) appeals to an algorithm that exploits the information that these components carry about external affairs and this system has been stabilized by some recruiting process.[12]

As a matter of fact, **?**, p. 213-16 accepts that representations extend well beyond the human mind and recognizes that some processes within plants or hormones probably qualify as representational. Nonetheless, he seems to oppose the idea that typically detectors are representations because he requires an additional constraint: according to him, a genuine representational system must possess what he calls 'robust outcome functions' i.e. effects that are produced in response to a range of different inputs and in a range of different relevant external conditions (see also **?**). Only detectors whose effects are robust in this sense can qualify as full-blown representations, and he argues that probably detectors often fail to satisfy this condition (see also **?**; **?**). As a response, I would like to make three remarks.

First, one could question the idea that robustness is *required* for the detector's semantic content to be explanatory. If the system often successfully produces a world-involving outcome by exploiting the information that internal components carry about external world affairs and this process has been designed or maintained by some stabilizing process, this should be enough for securing the explanatory bite of content attributions. How do moths manage to escape from bats? By possessing internal components that track ultrasounds, others that represent bats at a certain location and so on. I think ascribing content to them contributes to making sense of their successful behavior. Certainly, if the system were more robust it will probably be more reliable, moths will succeed more frequently and it would possess other features that we deem interesting. However, the question we are addressing is whether robustness should be *required* for content to be explanatory at all.[13] I think the reasons I provided for thinking that the detector's content contributes to en explana-

---

[12] This description does not rely on an explicit distinction between sender and receivers because Shea does not appeal to them (see Shea, **?**, p. 19).

[13] Note that even if members of kind K tend to possess property F, it does not follow that F is an important or explanatory property of K; F could be an accidental feature or a side-effect of a truly explanatory property, for example.

tion of behavior would hold even if typically representational systems were not robust.

Secondly, let us grant for the sake of the argument that robust effects are required. If 'robustness' is understood weakly enough (so weakly that I doubt Shea would accept it), most detectors would qualify. Neurons in the early auditory system of guinea pigs, for instance, change their sensitivity to sounds over time during sustained input to the neuron. This is probably an adaptive response to tune their sensitivity to the local sensory environment (Dean et al. 2005). Likewise, it has been suggested that the stimulus-response properties of AN2 interneurons change in a similar way (**?**, p. 10545). If all that is required for a mechanism to exhibit robustness is producing the same representation given this variety of inputs and conditions, many detectors would meet these criteria. Thus, a weak interpretation of 'robustness' is also compatible with the explanatory role of detectors.

Finally, there is a different way of showing that Shea's account is compatible with Strong Liberal Representationalism. **?** outs forward a pluralist proposal: he seeks to provide sufficient rather than necessary conditions for a state to qualify as a representation. Thus, the fact that he is interested in defining representations deriving from systems with robust outcome functions is still compatible with the existence of other kinds of genuine representations lacking this feature. Detectors could be one such case.

Summing up, I think Shea's model of how semantic properties explain behavior can also be used to argue that the detector's content is genuinely explanatory. I have presented two different models, which plausibly can accommodate the explanatory role of detectors and, consequently, support Strong Liberal Representationalism. Still, one could object that I chose the models that best suited my philosophical purposes. To address this concern, I will discuss a last proposal, which has a special feature: it has been provided by one of the most prominent critics of Liberal Representationalism. In particular, I will argue that according to **?**'s account of the explanatory role of semantic properties, the detector's content is explanatorily relevant.

2.3 Ramsey's approach

? argues that for a state to qualify as a genuine representation, it has to meet what he calls 'the job description challenge', which he defines as explaining 'how a physical state actually fulfills the role of representing in physical or computational process - accounting for the way something actually serves as a representation in a cognitive system' (**?**, p. xv). His own response to this challenge is that at least two kinds of states can be shown to play a representational role: Input-Output representations and Structural representations (henceforth, 'IO-' and 'S-representations'). Very roughly, S-representations are mental mechanisms that bear some form of structural resemblance to the entities they stand for, and which are employed as a 'model' for reasoning about the world. The paradigm example of S-representation is the cognitive map by means of which some organisms navigate the environment, although many other forms of structure-preserving processes can also play this role. Now, interestingly, **?** argued that detectors could actually be considered as a form of S-representations; for instance, the fact that an internal neuron can be in different states (active, inactive), which correspond to different world events (F, no-F), can be understood as a very simple form of structural mapping (see also **?**). Since S-representations uncontroversially meet the job description challenge, if Morgan is right and detectors are S-representations, then they should also qualify as genuine representations after all. Accordingly, Morgan's argument would provide a straightforward vindication of STRONG LIBERAL REPRESENTATIONALISM (even if he does not seem to endorse the latter).

Nonetheless, I think the representational status of detectors can also be defended by focusing on the other kind of states that, according to Ramsey, meet the job description challenge: IO-representations. Again, I think this notion is easier to grasp with an example that he provides (and which derives from **?**, p. 89). Suppose we try to build a machine for adding numbers. Numbers are abstract entities, so no physical mechanism can carry out this task by manipulating numbers. However, we can build a machine that possesses symbols that stand for numbers and the machine can then perform mathematical operations by manipulating these symbols. A picture of this relationship can be seen in figure **??**.

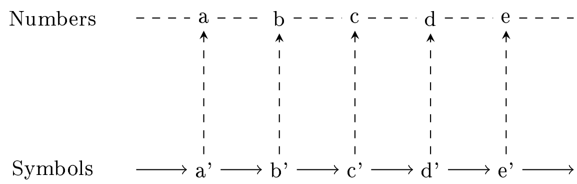

Fig. 2 Solid arrows stands for causal connections and dashed arrows for representational relations. Adapted from Ramsey (2007).

According to one way of understanding the Computational Theory of Mind, the same kind of scheme applies to mental processes. **?**, p. 69 argues that in order to explain "how an acoustic input that represents a certain public-language sentence winds up generating a representation of, say, a parse-tree for that sentence" we need to suppose that intermediate states are representations. Thus, certain data structures must be seen as representations because unless we suppose that the mechanism's sub-processes manipulate states that stand for other things we will not be able to understand how the mechanism performs its job (see also Milkowski, 2015, p. 717). These representational sub-routines must be in place for the mechanism to carry out the most sophisticated capacity that we try to explain. As Ramsey (2007, p. 74) argues:

> Serving as a representation of some feature of a target domain here amounts to serving as the sort of input or output required by a sub-

processor solving a problem related to that domain. The content of the representation is critical for this role because unless the symbol stands for the relevant computational argument or value, it is impossible to make sense of the sub-system as a computational sub-system doing its job.

My point is that this is exactly the role detectors play in the previous examples: We want to understand a distally-characterized capacity of crickets, namely predator avoidance. Unless we suppose that auditory detectors firing represent the presence of ultrasounds and activity in AN2 cells represents bats, we will not be able to understand how the insect manages to carry out this capacity. Thus, spike trains in auditory detectors and in AN2 cells should qualify as IO-representations. And, again, this reasoning seems to generalize to all detectors.

Further evidence for this interpretation is provided by Ramsey himself, since he briefly suggests that some kinds of detectors qualify IO-representations: "A theory about how the visual system extracts shape from shading is actually a theory about how we convert representations of shading into representations of shape" (**?**, p. 69). I think some states in early visual processing might qualify as detectors (see Ramsey, **?**, p. 120), so after all, his own notion of IO-representation can be used to vindicate STRONG LIBERAL REPRESENTATIONALISM.[14]

Summing up, I considered three different models of how semantic properties explain the system's behavior, Dretske's, Shea's and Ramsey's, and I argued that according to a reasonable interpretation of all of them detectors are genuinely explanatory. I suggested above, I think this result not only shows that the objection based on the explanatory idleness of representational attributions in the context of detectors is mistaken, but it also provides a positive reason for thinking that detectors are representations.

Crucially, note that as far as its explanatory status is concerned, detectors are not marginal in any respect: they fulfill this role as clearly as any other representation (although, of course, other representations might be more interesting for other reasons). Thus, these models not only support LIBERAL REPRESENTATIONALISM, but also STRONG LIBERAL REPRESENTATIONALISM.

[14] One might complain that this analysis mischaracterizes Ramsey's notion of IO-representations, since at some points he seems to suggest that for a process to count as an IO-representation, the system's output needs to qualify as a representation. This reply, however, is unsatisfactory for various reasons. First, on this reading, the only kinds of representations that do not presuppose the existence of any other sort of representation within the system are S-representations. Furthermore, it seems to be in tension with some of the claims he makes at other places (e.g. Ramsey, **?**, p. 77). Finally, the reasons adduced to support the status of IO-representations as genuine carriers of semantic content do not underpin this constraint. What seems to be crucial for the system's internal sub-processes to qualify as representational is that that the output be characterized distally. Adding that the output should itself qualify as a genuine representation seems unmotivated.

3 Guiding Research

I just argued that the detector's content can explain behavior. There is a second platitude about representations that should be true of detectors if we want them to possess semantic properties: this attribution should be useful in guiding research (?). We have to show that ascribing representational content to detectors suggests certain questions, whose answers significantly improve our understanding of certain phenomena. The hypothesis that they carry representational content should guide the investigation from an initial proposal to a deeper understanding of the mechanism. Let me briefly present some examples that show that this is indeed the case.

Once a representation is attributed, an important question that needs to be addressed concerns the properties of the representational vehicle. For instance, scientists knew that the firing of AN2 interneurons causes a steering response in crickets, but firing consists in different patterns of cellular activity, such as high-frequency bursts and isolated spikes. Which of them is AN2's representational vehicle? ?, p. 10542 asked that question and discovered that "bursts, but not isolated spikes, detect salient amplitude increases with high accuracy, (...) reliably signal the location of the stimulus (...) and reliably predict behavioral responses." Thus, the presupposition that detector activity carries semantic content led scientists to ask about the properties of their vehicles. Following that question, we learned that bursts, not isolated spikes constitute the relevant representational vehicle (Pollock, 2015).

Assuming that certain patterns of activity in AN2 constitute representations suggest some other interesting questions. For instance, is there any mechanism for dealing with noise? It has been suggested that the use of bursts rather than isolated spikes as representational vehicle can actually be understood as a strategy for filtering out noise (Lisman, 97; Swadlow and Gusev, 2001), but other mechanisms might in place. Would the sensitivity of AN2 interneurons change in different ambient conditions? It has been suggested that they do, like many other detectors (Smirnakis et al., 1997; Kvale and Schreiner, 2004; Dean et al., 2005; Marsat et al. 2006), but this issue still remains to be explored. In which conditions? How much can they change? An additional question concerns its representational content: What is activity in AN2 interneurons supposed to track? I have been assuming that it represents bats, but perhaps once we have a more complete computational description of the cricket's detection system, AN2 will be better described as representing some ultrasound properties. Similarly, a better ecological analysis might reveal the presence of other predators that are avoided through this system, so perhaps AN2 neurons also indicate the presence of some other predator. I think the answers to these questions are far from trivial and might provide important insights that significantly improve our understanding of this mechanism.

Furthermore, note that ascribing semantic content to detectors contributes to making certain useful generalizations. Consider the mechanism that insects employ to identify dead nest mates. Many eusocial insects use detection mechanisms to find out which members are dead and engage in sanitary measures (such as nest removal) to prevent the spreading of diseases or infections (?). Fire ants, for instance, remove dead nest mates when they sense that their bodies radiate certain cues like diglycerides and triglycerides (Choe et al., ?). On the perspective defended here, ants probably possess an internal representation of 'dead fellow', which triggers this prophylactic behavior of vital importance. Interestingly, note that different species probably track dead bodies by responding to different compounds (?). Furthermore, different ant species perform a diversity of behaviors: Leaf-cutter ants locate them in special refuse chambers (?), fire ants (*Solenopsis invicta*) scatter the corpses around the nest (?), the tiny *Strumigenys lopotyle* tends to create a tight ring of corpse fragments around its nest entrance (?), some myrmicine ants engage in cannibalism and ants of the species *T. lichtensteini* bury corpses with soil and nest material, to name just a few (?). However, the fact that all of them involve mechanisms for detecting (i.e. representing) dead mates enables us to classify them in the same category.[15] Likewise, by appealing to the ant's representational content we can classify their recognitional system along with the mechanisms of other hymenoptera species (?; ?). Consequently, the idea that ants represent corpses (rather than not representing anything or just representing proximal cues) makes certain generalizations possible that bridge over a variety of proximal stimuli and different kinds of actions.

I think this result generalizes to other detectors. For example, many species other than crickets (grasshoppers, moths, etc...) have evolved systems for detecting bats. So the strong liberal representationalist perspective suggests the following questions: do all other insects that are preyed by bats (such as beetles) possess similar mechanisms for recognizing their presence (?)? Are there some differences in the way they track bats? How stringent is the connection between bat detection and behavior? All these questions naturally follow from presupposing the representational nature of AN2 interneurons.

Therefore, I think the assumption that detectors are representations is actually useful: it suggests certain research questions in different areas that will eventually lead to a better understanding of the relevant mechanism.

4 Misrepresentation

The last platitude about representations that I would like to discuss is that their content can be false, i.e. there can be cases of misrepresentation. Although I think this is the easiest property to accommodate by SLR, it is important to avoid some possible misunderstandings.

To properly understand this desideratum, it is important to identify its motivations. On the one hand, as many people have pointed out, one of the more interesting and mysterious properties of representational states is that

[15] Of course, another aspect that unifies these mechanisms is that they trigger behaviors that deal with corpses. Given the fact that the signal's content probably depends on the behavior it elicits, one should expect a similarity of content to correlate with some similarity at the level of behavior (at least in simple cases).

they can be about non-existent facts. John can believe that Napoleon died in Waterloo, even if this fact (Napoleon dying in Waterloo) has never been the case. This is a feature that distinguishes representational phenomena from standard relations, since it is usually assumed that for A to stand in relation R to B, both A and B have to exist (Brentano, **?**; Jacob, **?**). This property also plays a key role in many arguments. For example, the existence of hallucinations lends support to a representationalist theory of perceptual experience precisely because these scenarios can be understood as involving experiences with false representational content (Fish, **?**, p. 66). Similarly, in the literature on naturalistic theories of content a key desideratum is to leave room for misrepresentation. For instance, a Naive Causal Theory according to which a state R represents S iff S causes R is widely rejected partly because on this account representational states cannot be false: if R represents S, S must cause R, so S has to exist (Stampe, **?**).

Now, from what I said in previous sections, it is not difficult to see that detectors can also misrepresent. The content of AN2 neurons firing is false iff there are no bats nearby. Likewise, the internal mechanism in ants used for detecting dead mates can also misrepresent: Choe and colleagues (**?**) found that the compounds that elicit necrophoresis are present in living and non-living organisms, the difference being that in the former there are additional 'chemical vital signs' that inhibit their detection. So a living ant that fails to produce these chemical vital signs will trigger a false representation of a corpse in other members of the colony, which will remove it from the nest. More generally, detectors misrepresent when they are set off, but the state they are supposed to detect is absent.[16]

Sometimes more stringent conditions for qualifying as a representation are assumed, which are related to the capacity to misrepresent, so some misunderstandings are possible in this context. Consider, for instance, a passage from a recent paper by Gładziejewski (2015: 76):

> (...) any truly representational system should have the ability to use representations to guide its actions even when the representational object is not present (see Chemero, 2009; Clark and Grush, 1999; Grush, 1997). In other words, we are not dealing with a truly representational system or mechanism unless the structures we want to treat as representations can be used off-line, outside of direct interactions with the representational object.

The property attributed in the second sentence significantly differs from the property ascribed in the first one. The fact that a state could have been false differs from the idea that it can be used off-line. The former is a reasonable requirement that derives from the general motivations concerning the notion of representation I sketched above, whereas the latter seems to unnecessarily

[16] If one thinks these mechanisms represent proximal features rather than distal events, misrepresentation is still possible. If AN2 neurons represent ultrasound waves rather than bats, then they misrepresent when detectors are directly stimulated by a clever neuroscientist using electrodes or by other cells that fail to work properly.

raise the bar. Unfortunately, sometimes there is some unclarity with respect to which of the two interpretation is put forward. For instance, Rowlands (2006: 157) claims:

> (...) whatever else a representation might be, it must be the sort of thing that can be used, by an organism, to guide its behavior in the absence of the feature of which it is a representation.

If that means that representations can obtain (and cause behavior) in the absence of the thing represented, it seems a plausible constraint and detectors possess it (since this amounts to the misrepresentation requirement). If it means something stronger –e.g. off-line reasoning or decoupability– it is unclear that there is anything in the notion of representation that commits us to this constraint.

Similarly, Milkowski (2013) argues:

> The account should also make misrepresentation possible; misrepresentation is still a kind of representation, from which information should be distinguished (...). The cognitive, rather than purely functional, role of misrepresentation must also be made clear. (...) Error should be system-detectable.

The claim that representational content can be false is very different from the idea that the system should be able to detect mistakes (cf. Bickhard, 1999, 2004; Milkowski, 2013; Gładziejewski, 2015: 76). Indeed, the ability to detect its own mistakes seems to require some sort of metarepresentational ability, and it is unclear why we should require metarepresentations in order to possess first-order representations. Of course, one might be interested in mechanisms that satisfy this additional constraint, but whereas the possibility of misrepresentations seems to be essential to identify something as a representation for the reasons outlined above, this not obviously true of the most stringent interpretations.

In conclusion, I think that a proper interpretation of the misrepresentation desideratum shows that detectors clearly fulfill it. Detectors can be false or unsatisfied, so they meet the three central conditions for qualifying as genuine representations.

5 Objections

So far I argued that SLR not only follows from some of our best naturalistic theories, but we also have independent reasons for thinking SLR is correct: ascribing representational content to detectors explains behavior (according to three different models), it is useful in guiding research and can account for misrepresentation. In this last section, I would like to address two objections.

First of all, one might object that a representational perspective does not add anything of value to a purely 'syntactic' account that merely appeals to the causal process that goes from stimuli to behavior (Schulte, **?**). To explain how the cricket managed to avoid a bat, we just need to appeal to the presence of a bat and its ultrasounds, the fact that they caused the activation of sound detectors, which caused the activation of ascending neurons AN2, which causes a steering reaction. Why do we need to appeal to the state's semantic content on the top of that? What does an interpretation in terms of content add to a purely causal account of behavior?

As a first reply, let me stress that in section **??** I extensively argued that according to three different models of how content explains behavior, the detector's semantic properties are explanatorily relevant. A purely 'syntactic' account misses certain relationships between internal states and environmental features that, among other things, help make sense of the organism's successful performance of a world-involving behavior (see section **??**). Nonetheless, let me make an additional point to illustrate why a pure causal account is probably insufficient.

In section **??** I focused on successful performance, but since in the last section I developed the idea that detectors can be false, we can now discuss how an attribution of false representational content can contribute to an explanation of unsuccessful behavior. For instance, consider pheromone traps, which are used to lure insects such as clothes moths into snares that will eventually kill them. Why do moths systematically move in the direction of the trap? Well, because some internal mechanism is representing the presence of a female moth that is willing to mate. The fact that internal states carry representational content about females helps us to understand why they behave as they do.

Crucially, note that in this case there is no alternative causal account available that can account for the way moths behave: what explains that a clothes moth gets attracted to a pheromone trap is something that is not present (females willing to mate), so one cannot explain this behavior by mentioning the fact that a female is causing it, because it is not. Likewise, merely asserting that clothes moths are built in such a way that pheromones attract them, is not even an explanation. More generally, I think the explanation of unsuccessful behavior is especially useful for illustrating the explanatory role of detectors because in this case there is no alternative 'syntactic' explanation, given that the properties appealed to are not instantiated.

One might reply that a causal approach can actually account for this behavior: we just need to add the claim that pheromones correlate with females willing to mate. But note that pheromones correlate with many others facts (nighttime, the presence of a male willing to mate nearby, a forest,...), so why are females singled out as the relevant property? Furthermore, there might actually be more pheromone traps than female moths sending these molecules, so it is not obvious that pheromones correlate better with females than traps. One move in the right direction is to insist that moths get attracted to pheromones because they have an internal component that has the evolutionary function of steering the moth toward a female when pheromones are sensed. This response, however, leads us to the next objection that I would like to consider.

According to this second objection, a representational perspective does not add anything of value to an account that appeals to causal processes, biological functions, correlations and mechanisms. Although a simple causal account that merely appeals to what is present cannot succeed (e.g. because of misrepresentation), if this account is supplemented with correlations, functions and the like, then adding an appeal to content is unnecessary. In the case of detectors, an appeal to representational content does not provide any additional, illuminating insight.

The problem with this suggestion is that it relies on the premise that an account in terms of causal processes, biological functions, correlations and the like is an alternative to an explanation in terms of content. However, this is precisely what naturalistic accounts such as teleosemantics deny. As I argued in section **??**, if teleosemantics (or some other naturalistic theory along the same lines) is correct, for R to be a representation with such and such content just is for R to be a state with certain functions, correlations and so on. Thus an explanation in terms of functions, mechanisms and the like is not an alternative to an explanation in terms of representations (this is argued at length in [author]).

Therefore, I doubt that any of these objections succeed. Any explanation that does not appeal to semantic content is either incomplete, or not a real alternative to a representationalist account. In either case, STRONG LIBERAL REPRESENTATIONALISM prevails.

6 Conclusion

In this essay I defended STRONG LIBERAL REPRESENTATIONALISM, i.e. the claim many detectors are genuine representations. On the one hand, SLR follows from some of our best naturalistic theories. On the other, the ascription of representational content to detectors satisfies the three desiderata stated in **??**: the detector's representational content explains behavior, this assumption is helpful in guiding research and misrepresentation is allowed. Of course, representations can differ in many respects: some of them provide more robust explanations than others, some are conscious, some exhibit a language-like structure, some have a non-conceptual content, and so on. Accordingly, for certain purposes some representations might be more interesting than others. However, with respect to their status as carriers of semantic content, detectors are as genuinely representational as a belief or a desire. Therefore, if the arguments of this paper are on the right tack, the conclusion should be clear: we should embrace the idea that detectors are genuine representational states. This is not a bullet we need to bite, but a pleasing consequence we need to savor.[17]

[17] In the last round of revisions for this journal, Ganson's (**?**) paper came to my attention. He defends a similar view, although he provides different arguments and does not discuss some of the issues that I develop in this paper (especially sections **??** and **??**). In any case, a discussion of his interesting paper will have to wait for another occasion.

References

Bechtel, W. (2016), 'Investigating Neural Representations: The Tale of Place Cells'. *Synthese* **193**(5), 1287–1321.

Bermudez, J. L. (2003), *Thinking Without Words*, Oxford University Press.

Bot, A. N. M., Currie, C. R., Hart, A. G. and Boomsma, J. J. (2001), 'Waste management in leafcutting ants', *Ethology Ecology and Evolution* **13**(3), 225–237.

Braddon-Mitchell and Jackson, F. (2007), *The Philosophy of Mind and Cognition*, Blackwell.

Brentano, F. (1874), *Psychology from an Empirical Standpoint*, Routledge

Burge, T. (2010), *The Origins of Objectivity*, Oxford University Press.

Butlin, P. (forthcoming), 'Representation and the active consumer', *Synthese* pp. 1–18.

Chemero, R. (2009), *Radical Embodied Cognitive Science*, Bradford Book.

Choe, D., Millar, J. G. and M. K. Rust. (2009), 'Chemical signals associated with life inhibit necrophoresis in Argentine ants', *Proc Natl Acad Sci U S A* **106** (20): 8251—8255.

Clark, A. (1997), *Being There. Putting Brain, Body, and World Together Again*, Bradford Book.

Cummins, R. (1975), 'Functional Analysis', *Journal of Philosophy* **72**, 741–765.

Cummins, R. (1991), 'The role of mental meaning in psychological explanation', *in* B. McLaughlin, ed., *Dretske and His Critics*, Basil Blackwell.

Dennett, D. (1991), 'Ways of establishing harmony'. In B. McLaughlin, ed., *Dretske and His Critics*, Basil Blackwell.

Diez, L., Lejeune, P. and Detrain, C. (2014), 'Keep the nest clean: Survival advantages of corpse removal in ants', *Biology Letters* **10**(7), 1–4.

Dretske, F. (1988), *Explaining Behavior. Reasons in a World of Causes*, The MIT Press.

Eliasmith, C. (2000), *How neurons mean: A neurocomputational theory of representational content*, Unpublished Dissertation, Washington University in St. Louis.

Fidgor, C. (2018), *Pieces of Mind: The Proper Domain of Psychological Predicates*, Oxford University Press.

Fish, W. (2010), *Philosophy of Perception: A Contemporary Introduction*, Routledge.

Fodor, J. (1986), 'Why Paramecia Don't Have Mental Representations', *Midwest Studies in Philosophy* **10**(1), 3–23.

Ganson, T. (2018), 'The Senses as Signalling Systems', *Australasian Journal of Philosophy,*, **96**(3), 519–531

Ganson, T. (2020),'A role for representations in inflexible behavior', *Biology and Philosophy,*, **35**, 1–18

Gładziejewski, P. (2015), 'Explaining cognitive phenomena with internal representations: A mechanistic perspective', *Studies in Logic, Grammar and Rhetoric* **40**(53), 63–90.

Gładziejewski, P. and Milkowski, M. (2017), 'Structural representations: causally relevant and different from detectors', *Biology and Philosophy* **32**(3), 337–355.

Godfrey-Smith, P. (1996), *Complexity and the Function of Mind in Nature*, Cambridge University Press.

Godfrey-Smith, P. (2009), *Darwinian Populations and Natural Selection*, Oxford University Press.

Godfrey-Smith, P. (2014), *Philosophy of Biology*, Princeton University Press.

Horgan, T. (1991), 'Actions, reasons, and the explanatory role of content', *in* B. McLaughlin, ed., *Dretske and His Critics*, Basil Blackwell.

Howard, D. F. and Tschinkel, W. R. (1976), 'Aspects of necrophoric behavior in the red imported fire ant, solenopsis invicta', *Behaviour* **56**(1/2), 157–180.

Hutto, D. and Myin, E. (2017), *Evolving Enactivism: Basic Minds Meet Content*, The MIT Press.

Hutto, D. and Myin, E. (2012), *Radicalizing Enactivism: Basic Minds without Content*, The MIT Press.

Jacob, P. (2013), 'Intentionality.', *Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), URL = <https://plato.stanford.edu/archives/win2019/entries/intentionality/>.

Kandel, E.; J. H. Schwartz; T. M. Jessel; S. A. Siegelbaum and A. J. Hudspeth. (2013), *Principles of Neural Science*, McGraw Hill.

Marsal, G. and Pollack, G. (2012), 'Bursting neurons and ultrasound avoidance in crickets', *Frontiers in Neuroscience* **6**(1), 1–9.

Martinez, M. (2013), 'Teleosemantics and indeterminacy', *Dialectica* **67**(4), 427–453.

Marsat G. and G. Pollack. (2013), 'A Behavioral Role for Feature Detection by Sensory Bursts', *The Journal of Neuroscience* **26**(41), 10542–10547.

Miller, L and A. Surllyke (2001), *How Some Insects Detect and Avoid Being Eaten by Bats: Tactics and Countertactics of Prey and Predator*, *Bioscience* **51**(7), 570–581.

Millikan, R. G. (1984), *Language, Thought and Other Biological Categories*, The MIT Press.

Millikan, R. G. (2009), 'Biosemantics'. In A. Beckermann, B. P. McLaughlin and S. Walter *The Oxford Handbook of Philosophy of Mind*, OUP.

Morgan, A. (2014), 'Representations gone mental', *Synthese* **191**(12), 213–244.

Morgan, A. (2018), 'Mindless accuracy: On the ubiquity of content in nature', *Synthese* **195**(12), 5403–5429.

Neander, K. (1995), 'Misrepresenting and Malfunctioning', *Philosophical Studies* **79**, 109–141.

Neander, K. (1995), 'Prunning the Tree of Life', *British Journal for the Philosophy of Science* **46**(1), 59–80.

Neander, K. (2017), *A Mark of The Mental*, MIT Press.

O'Malley, M. A. (2014), *Philosophy of Microbiology*, Cambridge University Press.

Papineau, D. (2003), 'Is representation rife?', *Ratio* **16**(2), 107–123.

Pollack, G. (2013), 'Neurobiology of acoustically mediated predator detection', *The Journal of Comparative Physiology* **201**, 99–109.

Price, C. (2001), *Functions in Mind: A Theory of Intentional Content*, Clarendon Press.

Ramsey, W. (2003), 'Are receptors representations?', *Journal of Experimental and Theoretical Artifical Intelligence* **15**(2), 125–141.

Ramsey, W. (2007), *Representation Reconsidered*, Cambridge University Press.

Rescorla, M. (2013), 'Millikan on honeybee navigation and communication', *in* D. Ryder; J. Kingsbury; K. Williford, ed., *Millikan and her Critics*, Blackwell.

Sakura, T., Namiki, S. and Kanzaki, R. (2014), 'Molecular and neural mechanisms of sex pheromone reception and processing in the silkmoth bombyx mori', *Frontiers in Physiology* **5**, 1–29.

Schulte, P. (2015), 'Perceptual representations: a teleosemantic answer to the breadth-of-application problem', *Biology and Philosophy* **30**(1), 119–136.

Schulte, P. (2018), 'Perceiving the world outside: How to solve the distality problem for informational teleosemantics', *Philosophical Quarterly* **68**(271), 349–369.

Schulte, P. (2019), 'Challenging Liberal Representationalism: A Reply to Artiga', *Dialectica*.

Shea, N. (2013), 'Naturalising representational content', *Philosophy Compass* **8**(5), 496–509.

Shea, N. (2018), *Representations in Cognitive Science*, OUP.

Sober, E. (1984), *The Nature of Selection*, MIT Press

Stampe, D. (2013), 'Toward a Causal Theory of Linguistic Representation', *in* P. French, H. K. Wettstein, and T. E. Uehling; K. Williford, ed., *Midwest Studies in Philosophy*, Minneapolis: University of Minnesota Press, **2**,, 42–63.

Sterelny, K. (1990), *The Representational Theory of Mind: An Introduction*, Oxford University Press.

Sterelny, K. (1995), 'Basic Minds', *Philosophical Perspectives* **9**,, 251–270.

Sterelny, K. (2003), *Thought in a Hostile World*, Blackwell Publishing.

Sun, Q. and Zhou, X. (2013), 'Corpse management in social insects', *International Journal of Biological Sciences* **9**, 313–321.

Yager, D. (2012), 'Predator detection and evasion by flying insects', *Current Opinion in Neurobiology* **22**, 201–207.

Wilson, E. O. (1971), *The Insect Societies*, Harvard University Press.