PhiMiSci
Philosophy and the Mind Sciences | Vol. 7 | 2026

Directive representation and the Job Description Challenge

Zina B. Ward[1,*]

Special Issue: *Representation in the Neurosciences and AI*, edited by Mark Sprevak & Francis Fallon

Abstract

Philosophical work on representation has largely focused on descriptive content, which concerns what the world is like. Neuroscientists and psychologists regularly make other sorts of content ascriptions, however, including to forward models, motor commands, and error signals. Here I focus on directive states – states that shape what a system does – in relation to the question: what makes a directive state a representation? I begin by recharacterizing Ramsey's (2007) Job Description Challenge (JDC) so that it applies to directive states. After arguing that existing answers to the JDC do not answer the "directive JDC," I develop my own account by asking: what is the *explanatory payoff* of appeal to directive representations? I argue that directive representations help explain how a system achieves a particular outcome in the face of varying internal or external conditions. To be a representation, then, a directive state must make a specific causal contribution to bringing about an outcome in a way that is decouplable from any particular behavior of the system. This account captures both what is distinctive about directive representation and what it has in common with descriptive representation. It also implies that we should be much more sparing in attributing content to directive states.

Keywords: Decouplability • Directive content • Directive representation • Job Description Challenge • Mental representation • Neural representation • Perceptual constancy

[1] Florida State University, United States
* Primary contact: zward@fsu.edu

Ward, Z. B. (2026). Directive representation and the Job Description Challenge. *Philosophy and the Mind Sciences, 7.* https://doi.org/10.33735/phimisci.2026.12005

1 Hegemony in the philosophy of representation

Philosophical work on mental and neural representation has been fueled by a collection of well-worn cases. Fly detectors (Fodor, 1992), edge detectors (Hubel and Wiesel, 1962), and place cells (Bechtel, 2016) have been used to argue for or against representationalism, to illustrate different views of mental content, and to puzzle through challenges related to indeterminacy and misrepresentation. But these familiar examples are strikingly homogeneous: fly detectors, edge detectors, and place cells purport to represent something about the environment. The literature on representation has subsisted on a diet of such examples because the paradigmatic representational state, for philosophers, is a state about *what the world is like*.

This focus on descriptive representation is understandable given philosophers' interest in how we access the world. But as long as philosophical theories of representation aim to account for its role in the mind and brain sciences, it is inadequate. Neuroscientists and psychologists regularly ascribe other sorts of content to mental and neural states, including forward models, motor commands, error signals, and internal tracking states. A number of philosophers have begun calling attention to these neglected species of representation (Millikan, 1995; Mandik, 1999, 2003; Thomson and Piccinini, 2018; Rusanen, Lappi, Pekkanen, and Kuokkanen, 2021; Rusanen, Lappi, Kuokkanen, and Pekkanen, 2021). In this paper I focus on directive states, which shape what a system does rather than convey information about the world. I consider directive states in relation to a fundamental philosophical question: what makes a state a representation? Existing philosophical work is mostly silent about when a directive state qualifies as representational and when it does not. Here I argue that a directive state is a representation if and only if it brings about an outcome in a way that is decouplable from the system's behavior. This view captures what is distinctive about directive representation while acknowledging that descriptive and directive representation are connected through the concept of decouplability.

I begin in Section 2 by introducing the "Job Description Challenge" (JDC): the task of characterizing the functional role that a state must play in a physical system in order to count as a representation (Ramsey, 2007). Because the Challenge has thus far been illustrated using only descriptive examples, Section 3 introduces the "directive JDC," which applies to directive states. Existing solutions to the Challenge are canvassed in Section 4. Having argued that none of them clearly apply to directive states, I develop my own answer to the directive JDC in Sections 5, 6, and 7. I argue that appeal to directive content has an explanatory payoff when a directive state can contribute to bringing about an outcome via different sorts of behaviors. A directive state is therefore a representation only when it is decouplable from particular behaviors of the system. I compare this proposal to a salient alternative from Price (2001) in Section 8 and address several objections to my view in Section 9. Section 10 concludes.

2 The Job Description Challenge

A physical state is a representation if and only if it plays a particular functional role in the system in which it is embedded (Haugeland, 1991; van Gelder, 1995; Ramsey, 2007; Cao, 2022). In *Representation Reconsidered* (2007), William Ramsey argues that, because representation is a functional notion, we have to separate two questions:

(1) What makes a physical state or structure function as a representation?

(2) Assuming that a physical state or structure *is* a representation, what gives it the content it has?[1]

Ramsey dubs the task of answering Question (1) "the Job Description Challenge" (JDC). The Challenge requires that we give an account of what it is

[1] Ramsey claims that the second question has received the lion's share of philosophers' attention, even though the first is "far more pressing" (Ramsey, 2007, p. 30). Others, however, have questioned whether the second question is really overlooked (Rupert, 2018).

for a state or structure to play a representational role in a physical system. We need a characterization of what representations *do* so that "we can then determine if a given state or structure fits the bill" (Ramsey, 2007, p. 25).[2] Ramsey elaborates:

> [W]e need to be told, in presumably computational, mechanical or causal/physical terms, just how the system employs representational structures. Principally, there needs to be some sort of account of just how the structure's possession of intentional content is (in some way) relevant to what it does in the cognitive system. After all, to be a representation, a state or structure must not only have content, but it must also be the case that this content is in some way pertinent to how it is used. We need, in other words, an account of how it actually serves as a representation in a physical system; of how it functions as a representation. (p. 27)

An answer to the JDC is an account of the functional role of representational states and structures.

Ramsey claims that representation is a cluster concept: there are different types of representations "that share various nominal features, but with no real defining essence" that unites them (p. 9). He entertains several different representational notions that purport to belong to the cluster, rejecting some and endorsing others. This pluralistic attitude toward concepts of representation stands in contrast to his monistic characterization of the JDC. The Challenge requires identifying *the* functional role of representational states: "we would like to know what different types of representations perhaps have in common, qua representation" (Ramsey, 2016, p. 6). Our multiple concepts of representation should all reflect representation's unitary functional role.

[2] Although Ramsey's approach to representation is realist, realism isn't a necessary presupposition of the JDC. The Challenge can be articulated against a backdrop of a number of attitudes toward representation. For instance, it can be understood as the task of specifying what properties a state must have in order to be *usefully regarded* as representational.

Ramsey's discussion makes clear that any answer to the JDC must satisfy several interconnected constraints. I will discuss two here. The first, which I'll call the *Pre-Theoretic Constraint*, insists that the answer should be responsive to our intuitive, everyday understanding of representation. Representational notions are already present in our non-scientific understanding of the world. There would be no point in scientists calling a mental state a representation unless the scientific notion of representation overlapped somewhat with our pre-theoretic conception. Just as labeling a biological entity a "pump" licenses certain assumptions, inferences, and expectations about it, so too calling a mental state a "representation" encourages us to think about it in a particular way (Ramsey, 2007, pp. 12, 28). Scientific concepts of representation need not reproduce every aspect of our pre-theoretic notion, but they also should not be completely disconnected from it. According to Ramsey, the things we intuitively call representations include propositional attitudes, memories, maps, books, hand signals, and compass needle positions (p. 20). An answer to the JDC should capture features present in these commonsense examples so that the mental states deemed representations by scientists are "recognizably representational in nature" (p. 13).

Closely related to the Pre-Theoretic Constraint is what I'll call the *Goldilocks Constraint*: an answer to the JDC must "stee[r] the course between the Scylla of putting forth conceptions of representation that are too strong (because central aspects of representation are left unexplained) and the Charybdis of positing conceptions of representation that are too weak (because representation is reduced to something non-representational, uninteresting and ubiquitous)" (Ramsey, 2007, p. 31).[3] The function of representational states and structures cannot be too easy to fulfill, on pain of trivializing the representational theory of mind. Nor should it be too demanding, blocking content ascriptions to mental states that are usefully regarded as representations.

[3] The Goldilocks Constraint can be seen as an extension of the Pre-Theoretic Constraint, since on our intuitive conception, representations aren't ubiquitous but they do exist. I separate the constraints for clarity of exposition but nothing hangs on their independence.

Ramsey is most concerned with the former danger. He thinks certain scientific concepts of representation are too liberal, threatening pan-representationalism. One such concept is the "receptor notion" of representation. A "receptor" or "detector" is a state or structure whose activation is reliably correlated with something in the environment. Scientists employing the receptor notion of representation assume that, "because a given neural or computational structure is regularly and reliably activated by some distal condition, it should be regarded as having the role of representing… that condition" (p. 119). Ramsey argues that detecting something in the environment does not yet make a state "recognizably representational." Lots of things reliably respond to specific stimuli. Our immune system, for example, consistently reacts to infections, and yet we would not say that it has the functional role of representing infections. Ramsey worries that treating receptors as representations turns the claim that cognition involves representational states into the "remarkably boring thesis that cognition is a process that involves states that are triggered by specific conditions" (p. 125). Receptors should therefore be seen only as "causal mediators" or "relay switches" (p. 126). This influential discussion of the receptor notion can be seen as expanding on Ramsey's Goldilocks Constraint: a satisfactory answer to the JDC must not allow a mere receptor to count as a representation (see also Orlandi, 2014).[4]

[4] Ramsey's JDC is not universally accepted. Rupert (2018) argues that representation, as a theoretical posit, need not be answerable to our ordinary notion. He also defends the idea that "detector-style representations" are genuine mental representations when embedded in a cognitive system. Others agree that Ramsey's dismissal of receptors is too quick (Sprevak, 2011; Morgan, 2014; Cao, 2022). I do not take a stand on this debate about receptors. My argument only denies representational status to the directive analogue of receptors, which I dub "triggers" in Section 3.

3 The directive JDC

The distinction between the directive and the descriptive is typically drawn with respect to content or representation.[5] As Butlin (2021) explains, for instance, "[r]epresentations with descriptive content purport to say how things are and can be true or false... Representations with directive content tell the systems that consume them what to do and can be satisfied or unsatisfied" (p. 2). The JDC concerns states that may or may not have content: their representational status is precisely what is at issue. I'll therefore broaden the distinction, drawing a contrast between descriptive states, which carry information about the world, and directive states, which shape what a system does.[6] Descriptive states include descriptive representations as well as states that carry information about the world non-representationally. Paradigmatic descriptive states are perceptual or sensory states. Directive states, meanwhile, encompass directive representations as well as states that shape what a system does in a non-representational fashion. They include desires and intentions, which are directive representations, as well as states whose representational status is less clear, including motor schemas (Pavese, 2020), biological drives (Millikan, 1984), and motoneuron activations (see below).

The JDC can be posed about either sort of state: what makes a descriptive state or a directive state function as a representation? Standard presentations of the Job Description Challenge have focused on descriptive states. Recall that the Pre-Theoretic Constraint holds that an answer to the JDC should hew to our pre-theoretic understanding of representation. Almost all of the everyday examples of representations that Ramsey cites possess descriptive rather than directive content. Memories, maps, books, compass needle positions – these represent something about the world,

[5] I take "descriptive" and "directive" to be interchangeable with "indicative" and "imperative," respectively, which are also widely used.

[6] These characterizations are not mutually exclusive. A state could carry information about the world while also shaping the behavior of the system in which it is embedded. I think this is as it should be, since it makes room for pushmi-pullyu representations, which have both directive and descriptive content (Millikan, 1995).

rather than something to be done. Out of all of Ramsey's examples, only "traffic signals and signs" seem to be directive in character (Ramsey, 2007, p. 20).

To show how the Pre-Theoretic Constraint applies to directive states, additional examples are needed. Let's start with Millikan's (1995) example of blueprints. Blueprints seem to be representations with directive content, at least when they are telling a builder where to lay a foundation or raise a wall. Likewise, furniture assembly instructions are commands for how to put something together. Or consider the footprints on the floor of an airport body scanner. These have directive content: they are about where one must position one's feet. Our ordinary notion of directive representation also encompasses entities and states in the natural world. For instance, flowers often contain special markings, many of which are ultraviolet and so not visible to the human eye, that tell insects where to land to collect nectar or pollen (see Figure 1). These intuitive examples show that our pre-theoretic notion of representation encompasses directive in addition to descriptive states. The Pre-Theoretic Constraint therefore applies to directive representation too: an account of what makes a directive state a representation must capture the features of blueprints, scanner footprints, and flower markings that (at least sometimes) make them representational.[7]

Ramsey's unpacking of the Goldilocks Constraint likewise centers descriptive states to the exclusion of directive ones. Recall that this constraint demands that representational attributions not be too liberal, including

[7] The examples in this paragraph are intuitive non-mental representations (Ramsey, 2007, p. 20). Ramsey observes that the non-mental representations picked out by our ordinary notion usually involve an interpreting mind (a builder interprets a blueprint; a traveler interprets printed footprints). To answer the JDC, Ramsey (2016) argues, is to "explain how something in the brain can function as a representation without something else in the brain functioning as a conscious representation-interpreter" (p. 4). Readers may have clearer intuitions about commonsense descriptive representations than directive representations. Perhaps our intuitive notion of representation applies most clearly to states about how the world is. Even if this is so, the fact that we ascribe content to desires and intentions suggests that our pre-theoretic understanding of representation countenances directive examples.

Figure 1: A spectral comparison of a yellow daylily flower under visible light (left) and ultraviolet light (right). Photograph by David Kennard used in accordance with CC-BY-SA 3.0 license.

by denying that receptors count as representations. Receptors are descriptive states or structures: they carry information about something in the environment. Ramsey does not say much about what a similar sort of excessive liberality in the domain of directives would look like, although he does briefly mention that a door stop does not represent the content "stop moving here" (Ramsey, 2007, p. 111).

I suggest that a directive analogue to the receptor notion of representation is a "trigger": a state or structure that reliably brings about a specific behavior. Triggers shape behavior but do not function as representations in the systems that contain them. Borrowing an example from Price (2001), we can take the heart as a paradigmatic example of a trigger. Price writes, "the squeezing motions produced by the heart will normally cause the blood to

bring it about that oxygen and nutrients are distributed around the body. But it would seem very unnatural to describe the heart as *telling* the blood to distribute oxygen and nutrients" (p. 137). The heart's contractions reliably trigger the distribution of oxygen and nutrients, but it would be strange to say that the contractions represent a command to distribute oxygen and nutrients. The heart's role in this process is not representational in nature. Or consider a spinal motoneuron that sends projections out of the spinal cord, innervating a skeletal muscle. When the motoneuron fires, it makes the muscle more likely to contract. The motoneuron's firing is a trigger, not a representation: it does not *represent* the command "contract muscle!" but merely causes muscular contraction. Thus, an answer to the JDC must not treat mere triggers as representations.

We've now seen that, although the examples standardly used to characterize the JDC neglect directive states, the Pre-Theoretic and Goldilocks Constraints apply equally well to directives. The task of the "directive JDC" is to provide an account of what it is for a directive state to play a representational role in a physical system. The "job description" of directive representations should capture what is representational about the intuitive examples discussed above and steer a middle course between the Charybdis of being too stingy with representational attributions to directive states and the Scylla of counting mere triggers as representations.

Let's now take a brief tour through existing answers to the JDC. Because they too are geared toward descriptive states, they have little hope of answering the directive JDC.[8]

4 The job description of representation: Existing proposals

At a high level of abstraction, there is near consensus on the job description of representations: representations are states that *stand (in) for something else* (Haugeland, 1991; van Gelder, 1995; Grush, 1997; Clark, 1997; Ramsey, 2007; Anderson and Rosenberg, 2008; Orlandi, 2014; cf. Cao and Warren, 2025). There is, however, no settled understanding of what it takes for one thing to stand in for another (Bechtel, 1998).

Some authors suggest that a state stands in for something else when it "tracks" that thing (Sterelny, 2003; Anderson and Rosenberg, 2008; Morgan and Piccinini, 2018). According to what Hacohen (2022) calls the "shared conception" of representation, "to be a representation of X" is to "have the function of tracking X," where X is a distal entity in the environment (p. 701). This line of thinking suggests an answer to the JDC: a state is a representation when it tracks, or has the function of tracking, something in the world. Putting aside substantive objections to this view, it is not applicable to the directive JDC. Directive representations need not track entities in the environment.

Another answer to the JDC with questionable applicability to directive states invokes surrogative reasoning. On this view, a state of a physical system stands in for something, and hence serves as a representation, when it is used by the system as a surrogate in reasoning about something else (Swoyer, 1991). For instance, a system that implements a model or simulation will contain surrogative elements that stand in for things in the target domain (Grush, 1997; Ramsey, 2007). Again, however, this proposal is geared toward descriptive states. One can imagine a mental simulation about, say, packing items into a box, where each element in the simulation acts as a surrogate for a real object on the table in front of me. But it is less clear what it would mean for a directive state to stand in for something

[8] In claiming that philosophers have overlooked the directive JDC, some might worry that I am neglecting "action-oriented" theories of representation (Clark, 1997; Grush, 1997; Mandik, 2005; Anderson and Rosenberg, 2008). Some of these theories are about pushmi-pullyu representations, rather than pure directive representations, which are my concern here (Clark, 1997). Others are action-oriented in that they argue that our ascription of content to a representational state should be responsive to how that state guides behavior (Grush, 1997; Anderson and Rosenberg, 2008). There are two reasons such accounts are mostly orthogonal to the current project. First, they remain concerned with sensory (read: descriptive) states and how they guide behavior. Grush's (1997) "emulators" are representations of the animal's body (see Chemero, 2009, ch. 3), while Anderson and Rosenberg (2008) treat representations as "stimulus-registrations." Second, they are primarily aimed at answering Ramsey's Question (2), not (1): they presume that a state is representational, and then provide an account of its content.

else in a surrogative reasoning process.[9]

Still another family of answers to the JDC claims that the marker of representational status is decouplability. It's important to tread carefully here, as there are more notions of decouplability floating around than philosophers realize (Table 1). One conception focuses on the relationship between a stimulus (or state of the world) and the behavior of a system: decoupling occurs when a system can behave in different ways in response to the same stimulus (or state of the world). According to some authors, the "presence of a representational mind starts surfacing in an organism when there is independence between what the organism does and what is given to it" (Orlandi, 2014, pp. 59–60). Even if this is a good litmus test for when a system as a whole is representational, it does not help us determine whether a particular state is a representation or not. Not every state within a system that is capable of representation is actually representational.

A related notion of decouplability involves the stimuli present to a system and a state within the system. Many authors argue that a state represents something only when the state can be tokened even in the thing's absence, or under variable presentations of it (Beckermann, 1988; Haugeland, 1991; Clark and Grush, 1999; Anderson and Rosenberg, 2008; Burge, 2010). For example, for my mental state to be a representation of a bagel, it must be possible for the state to be tokened even when there isn't a bagel in front of me (because I am merely imagining it, say), and when I'm shown a bagel under different lighting conditions, partly occluded, and so on. One reason this sort of decouplability is seen as crucial is that it makes room for the possibility of *mis*representation.

Orlandi (2014) is one of the authors who develops this sort of view in detail, arguing that "representations are performance-guiding states that can occur in the absence of what they represent" (p. 109). They bring decouplability to bear on the JDC, claiming that "the 'job description' of a representation is a performance guiding state which is independent from its causal and statistical antecedents present in the stimulus" (p. 31). Unsurprisingly, since Orlandi's book is about vision, this understanding of decouplability is geared toward descriptive states, not directive ones. Orlandi is interested in whether a state is decoupled from the stimulus or features of the environment that it purports to represent. Directive states are typically not candidates for this sort of decoupling.

Another influential conception of decouplability is found in Burge's (2010) discussion of perceptual constancy. Burge draws a sharp distinction between "perceptual representation" and "sensory information registration," which is non-representational (p. 316). Sensory states merely reflect the state of an individual's sensory receptors, whereas perceptual representations "concer[n] a receptor-independent reality" (p. 398). On Burge's view, the transformation from sensation to perception is brought about by perceptual constancies: capacities to "represent a given particular or attribute as the same despite significant variations in proximal stimulation" (p. 274). These include, for example, the capacity to represent an object's size as the same even as it moves closer or farther away. In addition to size constancy, animals' perceptual systems are capable of shape constancy, distance constancy, motion constancy, and color constancy (pp. 409–10). Because he holds that the operation of constancies enables perception, Burge closely associates representation and decouplability: a state is perceptual, and hence representational, only if it is decouplable from any particular stimulation of an organism's sensory receptors.[10] Again, however, this notion of decouplability is not applicable to non-descriptive states. Perceptual

[9] Bill Ramsey (personal communication) has suggested that surrogative reasoning involves an inference from the properties of a representational entity to the properties of a target, but that this inference can be either descriptive or prescriptive: it can be inferred that the target is *like* the representation, or that the target is *to be made like* the representation. In the latter case, the representation has directive content. This may be the best way of making sense of surrogative reasoning using directives.

[10] Burge (2010) does not claim outright that decouplability answers the JDC. Cautiously, he writes, "since they are not characterized independently of the notions representation and perception, one cannot use the notion perceptual constancy as an independent 'criterion' to determine when one has a case of genuine perceptual representation and when one has a case of non-representational sensory registration…Still, it seems to me that in a rough, non-criterial way, perceptual constancies are necessary as well as sufficient for…perceptual representation" (p. 413).

constancies are, after all, perceptual.

Hence, leading answers to the JDC, developed with descriptive representations in mind, do not address the directive JDC. Nevertheless, I think the family of answers that appeal to various notions of decouplability are on the right track. What we need is a notion of decoupling that applies to directive states.

5 Explanation and directive representation

Many authors argue that legitimate representational ascriptions have an explanatory payoff (Pylyshyn, 1984; Price, 2001; Burge, 2010; Orlandi, 2014; Shea, 2018; Cao, 2022). A philosophical theory of representation should be guided by the contribution content makes to explaining a system's behavior. As van Gelder (1995) writes,

> a useful criterion of representation – a reliable way of telling whether a system contains them or not – is to ask whether there is any explanatory utility in describing the system in representational terms. If you really can make substantially more sense of how a system works by concretely describing various identifiable parts or aspects of it as representations in the above sense, that is the best evidence you could have that the system really does contain representations. Conversely, if describing the system as representational lets you explain nothing over and above what you could explain before, why on earth suppose it to be so? (p. 352)

Although van Gelder is a skeptic about computationalism, the explanatory approach he articulates is popular even among representationalists. Ramsey (2007) too recasts the JDC as the task of determining when there is an "explanatory benefit in describing an internal element of a physical or computational process in representational terms" (p. 34).

I believe that philosophers who emphasize the connection between representation and decouplability are onto something important about the explanatory payoff of representation. Attributing content to a state – directive or descriptive – is explanatorily useful because it indicates that the state is involved in producing some sort of constancy in the face of variability (see especially Burge, 2010). It is often a significant achievement for a system to achieve invariance (whether in behavior, internal states, or something else) in spite of change (in a stimulus, the state of the world, or something else). This is why not all systems are representational, and why it is explanatorily substantive to assert that a state or entity has representational status – it indicates that the state or entity plays a role in generating invariance against a backdrop of variability.

In the domain of perception, states that are reliably elicited despite variability in how things appear are highly useful to the system that contains them. Systems with such descriptive states are not hostage to stimuli being presented in the same way across encounters; they can learn about a thing under different conditions. Representing a stimulus even when it is absent also enables planning, imagination, and counterfactual reasoning. Moreover, because decoupled perceptual states are useful to the system that contains them, identifying them allows us to better understand the system. Attributing content to descriptive representations helps explain how the system's complex behaviors arise.

In order to apply the insight that there is a link between representation and decoupling to directive states, we need to seek out different notions of constancy and variability. One such proposal has been developed within a teleosemantic framework by Carolyn Price (2001). Price too begins from the idea that representational attributions must earn their explanatory keep. She argues that we should not ascribe content to a device "unless we can see how the possession of that kind of content might help to explain the behaviour to which the device gives rise" (p. 138). To a teleologist, the relevant sort of explanation has a normative dimension: the ascribed content must explain how the behavior is *appropriate*. Price applies this explanatory approach to devices with imperative content, which she identifies with representations of goals: the point of an explanation that appeals to imperative content "will be to present the behaviour as appropriate, given the goal represented by the imperative device that caused it. In other words, the point of the

Table 1: **Table 1. A sampling of decouplability concepts. Note that these concepts are not in competition as long as different notions of decouplability are relevant to different types of representation or representation concepts. *= Sterelny and Price arguably invoke multiple notions of decouplability, not just those listed here.**

stimulus		state of the system	Burge (2010); Orlandi (2014)
stimulus		action	Sterelny (2003)*
state of the world	is decouplable	action	Martinez and Klein (2016)
controller	from ...	real environment	Grush (1997)
behavior		function of behavior	Price (2001)*
state of the system		behavior	Beckermann (1988)

explanation will be to make sense of that behaviour by showing that it was produced in order to achieve that goal" (p. 138). The question, then, is how exactly directive representations demonstrate the appropriateness of the behaviors they prompt, and what this tells us about when a directive state is a representation in the first place.

Chapter 6 of Price's (2001) book, in which she answers these questions most directly, can be interpreted in different ways. On one reading, Price claims that imperative content can be ascribed to a system when two conditions are met: (i) the cooperating mechanism (i.e., consumer of the representation) selects physical movements that the system will produce, and (ii) those movements serve different functions on different occasions (pp. 139–40). Price defends condition (ii) with the example of a person raising their arm. The raising of an arm has no single, stable function, so we cannot explain the appropriateness of someone raising their arm by appeal to the function it serves. Instead, we have to point to the goal that the raising of the person's arm served *on this occasion* and the fact that person represented that goal. Conversely, Price argues, the heart is not an intentional device because the movement of the blood has the same function on every occasion (namely, to distribute nutrients and oxygen). We can explain the movement of the blood by appeal to this function without invoking imperative content possessed by the heart.

On this reading of Price, the sort of decouplability relevant to representations with imperative content is decouplability between a behavior and its function (Table 1). In answer to the directive JDC, this interpretation of her view implies that a behavior must be decouplable from any particular function to warrant attributing representational status to the state(s) that prompt it.

The problem, however, is that this view makes representational status curiously independent of the internal workings of physical systems. For example, the view suggests that the raising of one's arm must be explained in terms of representational states simply because it has, say, a communicative function in the context of a wave to a friend but a protective function when one shields one's eyes from the sun. But it is odd that nothing about the states of the motor system which are causally involved in arm movement is relevant to the determination that they are representations. Or consider a simpler organism: a bacterium that has a flagellum. The flagellum whips around to move the bacterium from place to place. The movements of the flagellum serve different functions at different times, including escape from a dangerous environment, food seeking, movement toward a light source, and so on. If movements with variable functions must be explained by appeal to imperative content, then one has to say that the flagellar movements are representationally controlled, no matter how simple the bacterium's internal workings. But this can't be right. There could surely be systems that produce the same behavior in service of different functions,

but whose behaviors are not mediated by directive representations. Thus answering the directive JDC by appeal to decouplability between behaviors and their functions is misguided.

There is another interpretation of Price's view on which it's not just that the behavior brought about by an imperative device serves different functions on different occasions, but also that *different* states prompt the behavior on those different occasions. I will revisit this reading of Price in Section 8, once I have laid out my own view.

6 State-behavior decouplability

I've now argued that there is something right in the view that representation is linked to decouplability, and that this link is responsible for representation's explanatory payoff. But the sorts of decouplability considered thus far have been either tailored to descriptive states (Section 4) or implausibly detached from a system's internal workings (Section 5). To find a notion of decouplability suitable for the directive JDC, I'll now return to the question: What is the payoff of appealing to directive representations to explain a system's behavior? What explanatory work is done by appeal to states of a system that possess content about *what to do*?

My suggestion is that attributing directive content to the states of a system helps explain how the system is able to reliably bring about an outcome despite variability in the behavior required to achieve it. Some physical systems contain states that causally contribute to the production of certain outcomes by prompting different behaviors in different circumstances. Attributing content to such states explains how the system achieves something despite varying internal or external conditions: we say that the system contains a state that represents the outcome to be brought about.

For instance, it is explanatorily useful to talk about a state of the premotor cortex representing grasping because tokening that state prompts an organism to grasp in different ways (e.g., using different combinations of digits or starting from different hand positions). We can appeal to this premotor state when, for example, we seek to explain how a non-human primate is able to grasp objects of different shapes or textures. The premotor state prompts downstream neural activity that is capable of making compensatory adjustments for objects that are, say, slippery or awkwardly shaped. Ascribing directive content to the premotor cortex makes sense of the primate's outcome-directed yet flexible behavior. Or consider a person's intention to get the final ingredient they need for dinner, which they might accomplish by stopping by the grocery store, asking a neighbor, going to the farmer's market, or grabbing what's in their garden. Which of these actions they take might depend on the time of day, what else they have to do, how much cash they are carrying, and so on. This intention is a directive representation. It is explanatorily useful because it helps explain how the person is (or would be) able to get the last dinner ingredient across a range of possible circumstances. Note that, as these examples suggest, appeal to directive content can be explanatory at very different levels: the relevant outcomes, behaviors, and states can be specific and granular or general and abstract (see also Sections 7 and 9).

Just like descriptive states that are decouplable from any particular stimulus presentation, directive states that are decouplable from particular behaviors are highly useful to the system that possesses them. When the environment is variable, a system trying to achieve a particular goal might need to modify its behavior to get closer to the goal. It is beneficial for the system to contain a state that can prompt any number of those needed behaviors.

In general, then, my claim is that representational ascriptions to directive states help explain a system's resilience in bringing about target outcomes. This account of the explanatory value of directive representations requires decouplability between directive representations and behaviors of the system.[11] To be a representation, a directive state cannot reliably

[11] There are several other philosophers who have discussed similar notions of decouplability. One is Beckermann (1988), who argues that we can ascribe "wants or want-like states" to a system only when, "in a variety of different situations [the system] exhibits a behavior which is apt to bring about [a goal state] and the behavior is not the same in all situations" (p. 129). Although Beckermann is interested in personal-level wants, I extend this basic idea to subpersonal directive

prompt the same behavior in all circumstances. If it does, there is little point in attributing content to the state, since its contribution to behavior can be explained in simple causal terms. There is no need for an intentional explanation because a straightforward causal explanation will do.

7 The job description of directive representations

I'll now use decouplability between a state of a system and its behavior to formulate an answer to the directive JDC. To play a representational role, a directive state must make a *specific causal contribution* to bringing about an *outcome* in a way that is *decouplable* from any particular *behavior*. In other words, the job of directive representations is to bring about a target outcome in a way that is flexible. A directive state is representational if and only if it can causally contribute to the achievement of the outcome via different behaviors (Figure 2).[12] Consider again the example of a neural state in the premotor cortex whose tokening causes grasping (the outcome). Imagine that it can do so by prompting different sorts of movements (the behaviors) depending on the starting position of the hand, the nature of the object to be grasped, the presence of injured digits, and so on. On my

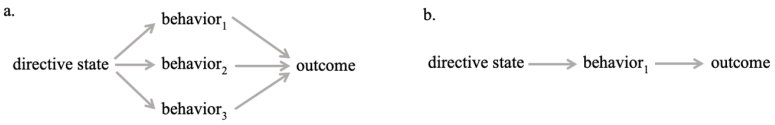

Figure 2: Figure 2. a. A directive representation. The directive state can prompt different possible behaviors that produce the outcome. b. A non-representational directive state. The state can bring about the outcome via one behavior only.

view, this premotor state counts as a representation because it can bring about grasping in a variety of ways.

Each of the italicized elements in the formulation above requires elaboration. First is the notion of a *specific causal contribution*, which invokes the concept of causal specificity from the literature on causal selection (Woodward, 2010). When a directive state makes a specific causal contribution to the production of an outcome, that means roughly that fine-grained changes to the state bring about fine-grained changes in the outcome. Requiring that the state's causal contribution be specific blocks the ascription of content to directive states that contribute to movement in a way that is not particularized to an outcome. For instance, there are states of the cerebellum that maintain posture and balance. Keeping one's balance is necessary for all sorts of actions. But those states of the cerebellum do not represent, say, "throw ball" just because they causally contribute to throwing. Their contribution is not causally specific with respect to that outcome.

Second, an *outcome* is the state of affairs that the directive state contributes to bringing about. Directive representations aim at states of affairs at different levels of abstraction. They might represent grasping an object, yawning, dancing, answering the phone, or buying a ticket. *Behaviors* also come at different levels of grain. Going to the ticket office, transacting online, or giving cash to the conductor are all ways of buying a train ticket; using different combinations of fingers or different arm trajectories are

states too. State-behavior decouplability can also be compared to what Sterelny (2003) calls "response breadth." An organism that is capable of response breadth has a "large menu of potential responses" when it registers something in the environment (p. 34). He introduces the term "decoupled representation" for "internal states that track aspects of our world, but which do not have the function of controlling particular behaviors" (p. 29). For Sterelny, however, decoupled representation is the analogue of our folk psychological notion of belief. His focus is on how beliefs, representational states with descriptive content, are not tied to specific behaviors. By contrast, I am arguing that states with directive content are decoupled from particular behaviors. Finally, the sort of decouplability that I am highlighting resembles what Butlin (2020) calls "one-many independence," a condition that he applies to consumers of representations, but which he ultimately rejects.

[12] My approach to answering the directive JDC is causal. See Section 9 for discussion of how similar ideas could be developed within a teleosemantic framework.

different behaviors that produce grasping.[13]

It may seem odd to entertain such a wide variety of outcomes and behaviors. Using these notions flexibly is necessary, however, because directive representations can be found at many different levels: we attribute directive content to premotor states and individual neurons as well as to personal-level intentions and desires. In order to accommodate these diverse directive representations, any answer to the directive JDC has to have similarly flexible elements. Moreover, the explanatory considerations discussed above support the ascription of content at varying levels. There are many different sorts of constancy of outcome that can be achieved in the face of different sorts of variability in what is required to bring it about. Whether a directive state is low-level and granular or high-level and abstract, ascribing content to it is explanatorily useful if it plays a particular role in the system in which it is embedded.

Finally, my characterization of the job description of directive representations requires *decouplability*, not actual decoupling (cf. Orlandi, 2014, p. 126). This gives the proposal a counterfactual dimension: what matters to whether a directive state is a representation is whether it *could* bring about an outcome in multiple ways, not whether it actually does. Imagine there is a directive state in your brain that, as a matter of fact, always brings about grasping in the same way. However, if you were to lose a finger, or if the path your arm takes were to be blocked by an obstacle, the state would bring about grasping by prompting different finger contractions or a different arm trajectory. Such a state is still reasonably considered a representation. It helps to explain why the system's achievement of grasping is resilient to possible perturbations, even if circumstances didn't happen to necessitate an adjustment.

This account of the job description of directive representations satisfies the constraints discussed in Section 3. The Pre-Theoretic Constraint requires that an answer to the JDC roughly capture what makes intuitive examples of directive representations – such as traffic signals, blueprints, and flower markings – genuinely representational. Consider a flashing yellow light at an intersection directing drivers to exercise caution. The light loosely embodies the sort of decouplability foregrounded here. It aims to bring about the state of affairs in which drivers are cautious. But there are multiple ways drivers can exercise caution: they can slow down, listen for other cars, hover their foot over the brake. State-behavior decouplability thus captures what makes the light a directive representation. Something similar might be said of flower markings that tell insects and birds where to land to gather pollen. Perhaps these markings bring about the arrival of pollinators by encouraging different behaviors (e.g., crawling, hovering, diving) that allow creatures to land on the relevant floral structures. If so, they are directive representations.[14]

The Goldilocks Constraint says that a characterization of the job description of representations should not lead to either excessively liberal or excessively sparing attributions of representational status to directive states. I argued above that this constraint is violated if triggers are treated as representations. My proposed answer to the directive JDC satisfies this constraint, for it implies that our earlier examples of triggers are not representational. The spinal cord is a trigger of muscle contraction rather than a representation of the command to contract because a motoneuron can only cause muscle contraction by firing action potentials, leading to the release of neurotransmitter at neuromuscular junctions. It does not have multiple behaviors at its disposal with which to bring about contraction. Likewise, the heart is not representational because although it contributes to the circulation of nutrients around the body, it can do so only in one way: by contracting and causing movement of the blood.[15]

[13] Because the grain of both behaviors and outcomes varies, something that is a behavior in one context can be an outcome in another. Grasping, for example, is sometimes an outcome and sometimes a behavior.

[14] Some might hold that these floral markings have descriptive rather than directive content: they represent *here is the pollen*, not *land here*. I think the best analysis of this case probably depends on biological details about the range of species which pollinate the flower, whether the markings always indicate the presence of pollen, how the pollinating species react to the markings, and so on.

[15] Another example of this sort is van Gelder's (1995) Watt governor. Most discussions of the Watt governor focus on whether the angle of its arms represent the current speed of the engine (e.g., Bechtel, 1998). But since arm angle controls the

This answer to the directive JDC captures both what is distinctive about directive representation and what directive and descriptive representation have in common. The account endorses the idea that what representations do is promote constancy in the face of variability. But it points out that different species of constancy and variability matter for directive states. The explanatory contribution of all representations is rooted in decouplability, but different sorts of decouplability are relevant for different types of representations.

8 Price (2001) redux

In Section 5 I argued against a potential answer to the directive JDC found in the work of Price (2001): a state has directive content when it brings about a behavior that serves different functions on different occasions. I showed that this view makes the representational status of a state implausibly insensitive to the internal workings of the system in which it is embedded. There is, however, another way to interpret Price. Consider the final clause in Price's contention that, "if we are to ascribe imperative content to a system, it must be the case that the system is able to trigger movements whose function varies from occasion to occasion, *depending on its causal origin*" (p. 139; my italics). Perhaps Price is here proposing a third condition on the possession of directive content, supplementing the two conditions discussed in Section 5: (iii) to be a directive representation, a state must be one of several states that can bring about a behavior, where the function of that behavior depends on which state prompted it. On this interpretation of Price, she agrees with my emphasis on decouplability between states and behaviors, but argues that it is a multiplicity of *states*, rather than of behaviors, that matters to directive representation. Such a view avoids my earlier objection, since it makes representational status dependent on which states are present within a system.

Consider the bee dance mechanism, which provides input to the bee's flight controller, which in turn produces physical movements, bringing the bee to a source of nectar. Price claims that the bee dance is an imperative device (that is, it possesses directive content). Under the current interpretation of Price, this is because all three conditions are met: (i) The flight controller produces different movements at different times. (ii) Those movements serve different functions on different occasions. The very same movement might serve the function "go to location L_1" or "go to location L_2" depending on external conditions like wind speed and direction (Price, 2001, pp. 140, 143).[16] (iii) Different bee dances (corresponding to different target locations) elicit a movement when it serves different functions. Thus, on the current reading of Price, bee dances are directive representations because they exhibit behavior-function decouplability (discussed in Section 5) in conjunction with a certain sort of state-behavior decouplability. But the latter is the mirror image of the decouplability that I have stressed: multiple states can cause the same behavior, rather than the same state causing multiple behaviors. My account gives the same verdict about bee dances but for a different reason. On my view, a bee dance is a directive representation because it can bring about an outcome – arrival at a specific location – by prompting whatever (flying) behavior is required to reach that location under prevailing conditions (wind speed and direction, etc.). Ascribing directive content to a bee dance is explanatorily useful because it highlights the dance's contribution to achieving a stable outcome.

My view has several advantages over the Pricean proposal. For one thing, Price's condition (ii) implies – implausibly, I think – that a behavior that always serves the same function cannot be controlled by a device with

opening of the throttle valve, we might also ask whether arm angle is a directive representation. I agree with van Gelder's skepticism about interpreting the Watt governor representationally, and my answer to the directive JDC substantiates this skepticism (cf. Chemero, 2000). The arm angles do not represent the valve opening because, although they are causally responsible for it, they can only bring it about in one way, namely by displacing a series of linked mechanical components between arms and valve.

[16] Although I will not press the objection here, one might wonder whether these movements actually serve a single function. Price claims that the bee dance system is "directed towards the satisfaction of a single need – the bees' need for nectar" but that "the specific function of each flight – to take the bee to a particular location – will vary from occasion to occasion" (p. 140).

directive content. Consider Price's example of a navigational system that "guides the organism back to its birthplace so that it can mate" (p. 140). The key difference between the navigational system and the bee dance system, according to Price, is that the function of the movements produced by the navigational system is always the same: to return the organism to its birthplace. So, according to Price, there is no "goal that needs to be represented" by the navigational system: imperative content does not earn its explanatory keep (p. 140). It seems to me, however, that a navigational system could be representational. Imagine that the system contains a state that, whenever it is activated, produces the behavior required at that particular moment for the organism to head in the direction of its birthplace, and different such behaviors are elicited in different conditions (e.g., swimming or flying, depending on whether the organism is in the air or in water). Ascribing directive content to the state would help explain how the organism manages to consistently return to its birthplace to mate.

Moreover, I explained in Section 5 that Price holds that an appeal to representational content has to show that a system's behavior is *appropriate*. This explanatory role for representation fails to motivate condition (iii) of the current Pricean proposal. It is not clear why the explanatory payoff of appeal to a directive state, and thus its representational status, depend on the existence of other, independent states that cause the same behavior.

Recall Price's account of how a directive state demonstrates the appropriateness of the behavior that it prompts: when the behavior can serve multiple functions, we must appeal to the representation of a goal that the behavior served *on this particular occasion* to show that it is appropriate. I have cast doubt on the idea that only behaviors whose function varies can be representationally controlled. But my view actually suggests an alternative way that directive representations explain the appropriateness of behavior. If a state is decouplable from any particular behavior, but coupled to a specific outcome, one can appeal to the content of the state to explain why the system's behavior was appropriate *under these conditions*. For instance: a particular bee flight was appropriate because the bee dance represented its target location L_t, and flying in such-and-such a way was needed to reach L_t under the prevailing winds. Thus, my view allows directive representations to explain why particular behaviors are appropriate without taking on the problematic commitments of the Pricean alternative floated here.[17]

9 Objections and replies

Before closing, I'll now try to head off a few further objections to my characterization of the job description of directive representations. A first worry relates to the claim that behaviors and outcomes have different levels of grain. If behaviors can be individuated extremely finely, it may seem that every directive state can be seen as representational. Consider our putatively non-representational example of the spinal motoneuron, which synapses on a skeletal muscle in a number of different places. Release of neurotransmitter at enough of these neuromuscular junctions will cause muscle contraction. So, one might say that the motoneuron has a number of different "ways" of causing muscle contraction, namely by causing neurotransmitter release at different neuromuscular junctions. Likewise, the heart can beat at different rates to pump blood around the body. One could claim that beating at 80 beats per minute (BPM) is a different behavior from beating at 90BPM.[18] These analyses would imply that the firing of the motoneuron and the beating of the heart are representational after all, running afoul of the Goldilocks Constraint.

[17] Indeed, at times it seems like Price has something quite similar to my account in mind. For instance, she claims that, in order to represent goals, a system must "be able to engage in a certain form of inference – one that involves the selection of a physical response that, in the circumstances, will realize one of a range of possible behaviors" (p. 144). This sounds like the sort of decouplability that I have stressed, so long as we interpret the "range of possible behaviors" as a set of behaviors that would bring about a particular outcome. Moreover, Price suggests that she is offering a teleological analogue to Beckermann's (1988) causal account, which is similar to mine (see footnote 11). Hence there seem to be at least three different answers to the directive JDC to be found in Price (2001).

[18] My thanks to Bill Ramsey for suggesting this example and the second possible reply to the objection, below.

The problem is that every directive state can be seen as decouplable from particular behaviors if there is no limit to how finely behaviors are individuated. By way of response, note first that other appeals to decouplability are vulnerable to the same worry (Green, 2023). Consider the claim that what makes a sensory state a representation is its decoupling from any particular stimulation of sensory receptors (Burge, 2010; Orlandi, 2014). It is possible to individuate stimuli extremely finely such that an object viewed from a particular angle and distance counts as a different stimulus than the object viewed from the same angle but 1mm farther away. Such individuation of stimuli would yield the conclusion that even early sensory states are decouplable and hence representational. This individuation problem is thus a difficult one for all accounts of representation that hinge on a notion of decouplability.

Let me sketch two possible directions of reply. The first is pragmatic: one might argue that the individuation of behaviors is fixed by the explanatory context. In many or most scientific contexts, it does not serve researchers' goals to individuate behaviors of the motoneuron according to the specific neuromuscular junctions at which neurotransmitter is released. Because distinctions that fine are usually irrelevant, we would not typically say that release of neurotransmitter at different sites constitutes different behaviors, and hence we can maintain that the motoneuron is not representational.[19] A potential downside of this reply is that it relativizes representational status to explanatory context. A state or entity is not a representation *tout court,* but only with respect to some context or other. It may be that there is some explanatory context in which differences in where neurotransmitter is released *are* salient, and so the motoneuron firing *would* count as a representation.

A different possible reply limits how behaviors are individuated by appeal to the consumer of the representational state. A key idea of teleosemantics is that a representation is something that is sent from a producer to a consumer, enabling the consumer to fulfill its proper function (Millikan, 1984). Even without committing to the details of this view, one might agree that being a representation requires being *treated as* a representation by a consumer system. That system could then be invoked as a backstop against the overly fine individuation of behaviors: two behaviors count as distinct only when they are treated as such by a consumer system. Of course, what is required for two behaviors to be distinguishable from the perspective of a consumer system would need to be spelled out in more detail. One might worry that this approach bottoms out in another individuation problem (substituting the individuation of proper functions, say, in place of the individuation of behaviors) or that it risks circularity (invariant consumer responses potentially being invoked both to individuate behaviors and as a requirement for decouplability). It remains to be seen whether the reply could be developed in way that avoids these potential pitfalls.

A second possible objection to my answer to the directive JDC is that it seems to be in direct conflict with the popular view that a mental representation has more directive content the *tighter* its connection to behavior (Lewis, 1969; Huttegger, 2007; Zollman, 2011; Martínez and Klein, 2016; Butlin, 2021). This idea is common among philosophers who adopt teleosemantic or information-theoretic theories of representation. Proponents of such views often think of representations as signals in signaling games. A signaling game involves a sender who observes the state of the world and then transmits a signal to a receiver, who acts on the basis of the signal. Following Lewis (1969), a number of authors have suggested that a signal has imperative content when the receiver always acts in the same way upon receiving it. By contrast, if the receiver "deliberates," or acts in different

[19] Burge (2010) might be interpreted as adopting this pragmatic approach to the individuation problem. He thinks the operation of perceptual constancies confer representational status on perceptual states. Perceptual constancies are "capacities systematically to represent a given particular or attribute as the same despite significant variations in proximal stimulation" (p. 274). But how "significant" is significant enough? How should proximal stimulations be individuated? Burge concedes that "there are surely borderline cases" between representational and non-representational systems (p. 410). But he seems to think that a discontinuity is introduced by a certain explanatory focus: "Described from the point of view of anatomy and physiology, there is a continuum between an amoeba's sensitivity to light and human vision. Described from the point of view of explanations of the visual systems in mammals and other relatively complex animals, perceptual representation is a distinctive kind of psychological state" (p. 319).

ways in different situations, then the signal has indicative content (Huttegger, 2007). This approach suggests that it is coupling between states and behaviors, rather than decoupling, that makes for directive representation.

However, the aforementioned work is not aimed at characterizing the job description of representations (Ramsey, 2016). These authors simply assume that a particular state or signal is representational, and then ask whether (or how much of) that signal's content is imperative versus indicative. Furthermore, characterizing the job description of directive representations in terms of coupling between states and behaviors would lead to radical over-ascription of representational status to directive states, violating the Goldilocks Constraint. Plenty of non-representational systems exhibit such coupling. Contractions of the heart are coupled with the blood's circulation of nutrients around the body. And yet the heart (a sender) is not telling the blood (a receiver) to distribute nutrients, even though every time it pumps (a signal), that's just what the blood does.

Instead of state-*behavior* coupling, we might instead take at least some teleosemanticists to be claiming that directive representations are coupled to distal states of affairs or abstractly characterized behaviors. Perhaps these target state of affairs or abstract behaviors can be realized by different proximal behaviors in different circumstances. If we reinterpret the point in this way, far from being an objection, it is effectively a teleosemantic analogue of my (causal) answer to the directive JDC. Instead of claiming that a directive state has to make a specific causal contribution to bringing about a particular outcome in order to be a directive representation, the teleosemanticist can claim that the directive state needs to *have the function of* bringing about such an outcome.

A final objection to my answer to the directive JDC is that it delegitimizes many of the content ascriptions made by psychologists and neuroscientists. On my view, unless a directive state plays a causally specific role in bringing about an outcome *and* it can do so by prompting different behaviors, the state is not properly called a representation.[20] One might object that setting such a high bar for directive representation is out of step with scientific practice. Researchers frequently talk about motor commands, goals, and other directive content being represented by states and structures that are not capable of eliciting behavioral flexibility.

It is important to keep in mind the general state of play with respect to representation-talk in psychology and neuroscience. Philosophers have long lamented that scientists throw the term "represent" around too casually. Even Cao (2022), who is skeptical of such philosophical critiques, concedes that "[s]ometimes representation talk is over-used, perhaps as mere shorthand for causal intermediaries that are not necessarily content-involving" (p. 22). There are neuroscientists who agree that representation is frequently invoked unnecessarily (Krakauer, 2022). Recent empirical research also shows that scientists are often uncertain about when it is reasonable to characterize brain activity representationally (Vilarroya, 2017; Favela and Machery, 2023). These studies focus primarily on content ascription to sensory states; I suspect that there is even greater uncertainty in other domains.

At present, then, there is a good deal of confusion about representation in the mind and brain sciences. In light of this confusion, it is reasonable to insist that directive representation is harder to come by than neuroscientists and psychologists tend to think. Just as Ramsey (2007) acknowledges that his rejection of the receptor notion of representation is "fairly radical," I grant that my answer to the directive JDC is quite revisionary (p. 148). But this a strength of the view, not a shortcoming. Scientists are in need of clear criteria for when it is appropriate to ascribe content to a directive state, and when doing so is explanatorily gratuitous.

10 Against descriptive hegemony

By applying the Job Description Challenge to directive states, this paper joins a growing chorus of philosophers calling attention to neglected

[20] This is in line with the view that the use of directive representations requires more cognitive sophistication than the use of descriptive representations (Price, 2001; Sterelny, 2003; cf. Mandik, 2003). As Price (2001) puts it, "the ability to represent goals is a more sophisticated capacity than the ability to signal the occurrence of some state of affairs" (p. 125).

species of representation. Recent years have seen increasing interest in how directive content differs from descriptive content (Millikan, 1984, 1995; Huttegger, 2007; Zollman, 2011; Artiga, 2014; Shea, 2018; Butlin, 2021). Mandik (2003), for instance, claims that it is unlikely that perceptual states get their content in the same way as all other representational states. He models the evolutionary emergence of different sorts of representations and encourages other philosophers to likewise look beyond perception (Mandik, 1999, 2003, 2005). Rusanen,Lappi, Kuokkanen, and Pekkanen (2021) examine the "goal-based representations" used in reinforcement learning, claiming that we lack a philosophical account of rewards and value representations (p. 14030). In another paper, they argue that philosophers need to "let go of the assumption that only states that track external environments count as representation" (Rusanen, Lappi, Pekkanen, and Kuokkanen, 2021, p. 1053). Figdor (2021) goes further, arguing that there are intentional states that are not representations at all. On her view, error signals, which are evaluative states, do not fit the representational mold.

A number of questions arise once we extend our attention beyond descriptive states. I have shown that we must recharacterize the JDC for directive states and formulate a new answer. One might wonder whether the descriptive and directive JDCs are exhaustive, or whether we will we need a catalogue of JDCs for different sorts of states. Representational attributions are often defended by invoking the idea that certain tasks are representation-hungry: they cry out for explanation in terms of internal representational states (Clark and Toribio, 1994; cf. Zahnoun, 2021). Are such tasks always hungry for descriptive representations, or are there tasks that seem to require directive representations too (cf. Keijzer, 1998)? Classic puzzles about representation, such as those related to indeterminacy and misrepresentation, have also been seen through an exclusively descriptive lens. What are the directive analogues of indeterminacy or misrepresentation, and how should they constrain our theories of content?[21]

In thrall to perception, philosophers have overlooked such questions. We would do well to further extend our thinking beyond the domain of the descriptive, heeding Millikan's (1984) prescient warning: "intentionality is not of a piece" (p. 86).

[21] Ward (2023) discusses what may be a real-life case of indeterminacy in the motor domain.

Acknowledgements

For valuable feedback and conversations, I would like to thank Sara Aronowitz, John Bickle, Zoe Drayson, Francis Fallon, Seth Goldwasser, Muhammad Ali Khalidi, John Krakauer, Edouard Machery, Corey Maley, Joe McCaffrey, Shivam Patel, Bill Ramsey, Tomás Ryan, Eden Sayed, Timothy Schroeder, and other members of audiences at the 2022 meeting of the Deep South Philosophy and Neuroscience Workgroup and the 2023 meeting of the Society for Philosophy and Psychology. I'm also grateful to two very helpful reviewers for this journal.

References

Anderson, M. L., & Rosenberg, G. (2008). Content and action: The guidance theory of representation. *The Journal of Mind and Behavior*, *29*(1/2), 55–86.

Artiga, M. (2014). Teleosemantics and pushmi-pullyu representations. *Erkenntnis*, *79*(3), 1–22. https://doi.org/10.1007/s10670-013-9517-5

Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist's challenge in cognitive science. *Cognitive Science*, *22*(3), 295–317. https://doi.org/10.1016/S0364-0213(99)80042-1

Bechtel, W. (2016). Investigating neural representations: The tale of place cells. *Synthese*, *193*(5), 1287–1321. https://doi.org/10.1007/s11229-014-0480-8

Beckermann, A. (1988). Why tropistic systems are not genuine intentional systems. *Erkenntnis*, *29*(1), 125–142.

Burge, T. (2010). *Origins of objectivity* (1st). Oxford University Press.

Butlin, P. (2020). Representation and the active consumer. *Synthese*, *197*(10), 4533–4550. https://doi.org/10.1007/s11229-018-01941-9

Butlin, P. (2021). Directive content. *Pacific Philosophical Quarterly*, *102*(1), 2–26. https://doi.org/10.1111/papq.12328

Cao, R. (2022). Putting representations to use. *Synthese*, *200*(2). https://doi.org/10.1007/s11229-022-03522-3

Cao, R., & Warren, J. (2025). Mental representation, "standing-in-for", and internal models. *Philosophical Psychology*, *38*(2), 379–396. https://doi.org/10.1080/09515089.2023.2207594

Chemero, A. (2000). Anti-representationalism and the dynamical stance. *Philosophy of Science*, *67*(4), 625–647.

Chemero, A. (2009). *Radical embodied cognitive science*. MIT Press.

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. MIT Press.

Clark, A., & Grush, R. (1999). Towards a cognitive robotics. *Adaptive Behavior*, *7*(1), 5–16.

Clark, A., & Toribio, J. (1994). Doing without representing? *Synthese*, *101*(3), 401–431. https://doi.org/10.1007/BF01063896

Favela, L. H., & Machery, E. (2023). Investigating the concept of representation in the neural and psychological sciences. *Frontiers in Psychology*, *14*. https://doi.org/10.3389/fpsyg.2023.1165622

Figdor, C. (2021). Shannon + Friston = content: Intentionality in predictive signaling systems. *Synthese*, *199*(1), 2793–2816. https://doi.org/10.1007/s11229-020-02912-9

Fodor, J. A. (1992). *A theory of content and other essays*. The MIT Press. https://doi.org/10.7551/mitpress/6765.001.0001

Green, E. J. (2023). Perceptual constancy and perceptual representation. *Analytic Philosophy*, *65*(4), 473–513. https://doi.org/10.1111/phib.12293

Grush, R. (1997). The architecture of representation. *Philosophical Psychology*, *10*(1), 5–23. https://doi.org/10.1080/09515089708573201

Hacohen, O. (2022). What are neural representations? A Cummins functions approach. *Philosophy of Science*, *89*(4), 701–720. https://doi.org/10.1017/psa.2022.6

Haugeland, J. (1991). Representational genera. In J. Haugeland (Ed.), *Having thought: Essays in the metaphysics of mind* (pp. 171–206). Harvard.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106–154.

Huttegger, S. M. (2007). Evolutionary explanations of indicatives and imperatives. *Erkenntnis*, *66*(3), 409–436. https://doi.org/10.1007/s10670-006-9022-1

Keijzer, F. A. (1998). Doing without representations which specify what to do. *Philosophical Psychology*, *11*(3), 269–302. https://doi.org/10.1080/09515089808573262

Krakauer, J. (2022). But is it thinking? The philosophy of representation meets systems neuroscience. *Studies in History and Philosophy of Science*, *92*, 267–269. https://doi.org/10.1016/j.shpsa.2021.05.014

Lewis, D. K. (1969). *Convention: A philosophical study*. Wiley-Blackwell.

Mandik, P. (1999). Qualia, space, and control. *Philosophical Psychology*, *12*(1), 47–60. https://doi.org/10.1080/095150899105927

Mandik, P. (2003). Varieties of representation in evolved and embodied neural networks. *Biology and Philosophy*, *18*(1), 95–130. https://doi.org/10.1023/A:1023336924671

Mandik, P. (2005). Action-oriented representation. In A. Brook & K. Akins (Eds.), *Cognition and the brain: The philosophy and neuroscience movement* (pp. 284–305). Cambridge University Press.

Martínez, M., & Klein, C. (2016). Pain signals are predominantly imperative. *Biology and Philosophy*, *31*(2), 283–298. https://doi.org/10.1007/s10539-015-9514-y

Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT Press.

Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, *9*, 185–200. https://doi.org/10.2307/2214217

Morgan, A. (2014). Representations gone mental. *Synthese*, *191*(2), 213–244. https://doi.org/10.1007/s11229-013-0328-7

Morgan, A., & Piccinini, G. (2018). Towards a cognitive neuroscience of intentionality. *Minds and Machines*, *28*(1), 119–139. https://doi.org/10.1007/s11023-017-9437-2

Orlandi, N. (2014). *The innocent eye: Why vision is not a cognitive process*. Oxford University Press.

Pavese, C. (2020). Practical representation. In E. Fridland & C. Pavese (Eds.), *The routledge handbook of philosophy of skill and expertise* (pp. 226–244). Routledge.

Price, C. (2001). *Functions in mind: A theory of intentional content* (1st). Clarendon Press.

Pylyshyn, Z. W. (1984). *Computation and cognition: Toward a foundation for cognitive science*. The MIT Press. https://doi.org/10.7551/mitpress/2004.001.0001

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

Ramsey, W. (2016). Untangling two questions about mental representation [Special Issue: Explaining Representation]. *New Ideas in Psychology*, *40*, 3–12. https://doi.org/10.1016/j.newideapsych.2015.01.004

Rupert, R. D. (2018). Representation and mental representation. *Philosophical Explorations*, *21*(2), 204–225. https://doi.org/10.1080/13869795.2018.1477979

Rusanen, A.-M., Lappi, O., Kuokkanen, J., & Pekkanen, J. (2021). Action control, forward models and expected rewards: Representations in reinforcement learning. *Synthese*, *199*(5), 14017–14033. https://doi.org/10.1007/s11229-021-03408-w

Rusanen, A.-M., Lappi, O., Pekkanen, J., & Kuokkanen, J. (2021). From fly detectors to action control: Representations in reinforcement learning. *Philosophy of Science*, *88*(5), 1045–1054. https://doi.org/10.1086/715513

Shea, N. (2018). *Representation in cognitive science*. Oxford University Press. https://doi.org/10.1093/oso/9780198812883.001.0001

Sprevak, M. (2011). Review of William M. Ramsey *Representation Reconsidered. British Journal for the Philosophy of Science*, *62*(3), 669–675. https://doi.org/10.1093/bjps/axr022

Sterelny, K. (2003). *Thought in a hostile world: The evolution of human cognition* (1st). Wiley-Blackwell.

Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, *87*(3), 449–508. https://doi.org/10.1007/BF00499820

Thomson, E., & Piccinini, G. (2018). Neural representations observed. *Minds and Machines: Journal for Artificial Intelligence, Philosophy and Cognitive Science*, *28*(1), 191–235. https://doi.org/10.1007/s11023-018-9459-4

van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy*, *92*(7), 345–381.

Vilarroya, O. (2017). Neural representation. A survey-based analysis of the notion. *Frontiers in Psychology*, *8*. https://doi.org/10.3389/fpsyg.2017.01458

Ward, Z. B. (2023). Muscles or movements? Representation in the nascent brain sciences. *Journal of the History of Biology*, *56*(1), 5–34. https://doi.org/10.1007/s10739-023-09703-1

Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology and Philosophy*, *25*(3), 287–318. https://doi.org/10.1007/s10539-010-9200-z

Zahnoun, F. (2021). On representation hungry cognition (and why we should stop feeding it). *Synthese*, *198*(1), 267–284. https://doi.org/10.1007/s11229-019-02277-8

Zollman, K. J. S. (2011). Separating directives and assertions using simple signaling games. *The Journal of Philosophy*, *108*(3), 158–169.