

Representation Re-construed: Construal-based Norms for Ascribing Natural Representations

Abstract: Many philosophers worry that cognitive scientists apply the concept REPRESENTATION too liberally. For example, William Ramsey argues that scientists often ascribe natural representations according to the “receptor notion,” a causal account with absurd consequences. I rehabilitate the receptor notion by augmenting it with a background condition: that natural representations are ascribed only to systems construed as organisms. This Organism-Receptor account rationalizes our existing conceptual practice, including the fact that scientists in fact reject Ramsey’s absurd consequences. The Organism-Receptor account raises some worrying questions, but as a more faithful characterization of scientific practice it is a better guide to conceptual reform.

Abstract: 100 words

Total: 4,995 words

1. Introduction. There is a common complaint among philosophers that scientists use the word “representation” too liberally. Representation is often contrasted with indication: representation is a distinction achieved by maps, linguistic performances, and thoughts, whereas indication is a less-demanding state achieved by thermostats, which indicate ambient temperature, and refrigerator lights, which indicate whether the door is open (Dretske 1981; Cummins and Poirier 2004). However, cognitive scientists often ascribe representations when it seems that mere indication is all that is called for. We commonly say that hidden layers in a neural network represent concepts, or that neurons in V1 represent visual edges, because they reliably respond differently to the circumstances they are said to represent (Ramsey 2007, 119–20; cf. Hubel and Wiesel 1962). But these “representations” are thin-blooded compared to paradigmatic conventional representations. For example, they cannot be invoked in the absence of an appropriate stimulus. So are cognitive scientists conceptually confused? Do they exaggerate their claims? And if the natural representations posited by cognitive scientists aren’t genuine representations, is the cognitive revolution dead?

William Ramsey provides an excellent book-length exploration of these worries, articulating a qualified pessimism about their answers:

...we have accounts that are characterized as “representational,” but where the structures and states called representations are actually doing something else. This has led to some important misconceptions about the status of representationalism, the nature of cognitive science and the direction in which it is headed. (2007, 3)

Ramsey describes the “job description challenge”: to give an account of the distinctive properties of representations in virtue of which appealing to them serves a special

explanatory role. If the job description challenge can be met, then we can formulate a plan for conceptual reform.

I undertake Ramsey's challenge, but with a metadiscursive twist: I describe the Organism-Receptor account, which articulates conditions for ascribing representations, in virtue of which such ascriptions achieve a special explanatory purpose. The account is merely suggestive about the properties that distinguish first-order representational states from non-representational states; it says more about the mental state of the ascriber than about the representation-bearing system. However, the Organism-Receptor account provides a more adequate characterization of scientists' practice than Ramsey's.

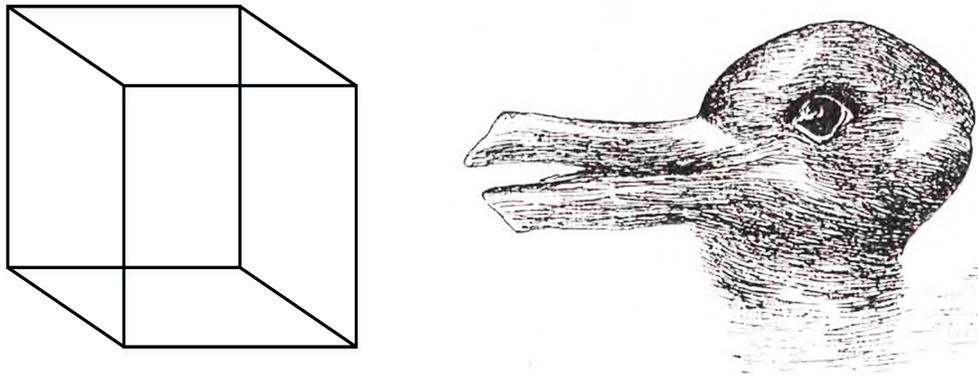
My main aim in this paper is to push back against pessimistic evaluations of the existing practice of representation-ascription in cognitive science, like Ramsey's. I will focus on Ramsey's critique of the "receptor notion," a flawed causal theory of representation that he attributes to some cognitive scientists. Ramsey argues that the receptor notion has absurd consequences, although scientists do not accept them. By augmenting the receptor notion with a construal-based background condition, I can explain why scientists do not draw these absurd conclusions. Whereas Ramsey's pessimistic account of scientists' practice of ascribing representations finds it wanting and is extensionally inadequate, mine rationalizes our extant conceptual practice (though that practice is not beyond criticism). I conclude that my apologetic account is a more charitable and adequate interpretation of existing scientific practice than Ramsey's.

2. Ramsey on the "Receptor Notion." Ramsey argues that natural representations in cognitive science are often ascribed according to the "receptor notion," a crude causal theory of representation. According to the receptor notion, a state s represents a state of affairs p if s is regularly and reliably caused by p (2007, 119).

Ramsey claims that the receptor notion is what justifies the ascription of representations to cells in V1 that detect visual edges, cells in frog cortex that detect flies, and the mechanisms in Venus flytraps that cause their “jaws” to close (119–23). Ramsey argues that this receptor notion is too liberal to be useful to scientists. For example, it is susceptible to the “disjunction problem” (Fodor 1987): since frog neurons respond reliably to visual stimulation by flies *or* (say) BBs, we should say that the content of the representation is *fly-or-BB*, rather than *fly*. Likewise, Venus flytraps represent objects in a particular range of sizes rather than *edible insects*, and the human concept GOAT represents *goats-or-weird-looking-sheep*. Such disjunctive content-ascriptions are usually considered absurd. Absent a clever fix, we must embrace unwieldy, disjunctive contents for representations or we must reject the receptor notion (Ramsey, 129).

Dretske’s (1988) teleofunctional theory of representation is a sophisticated twist on the receptor notion that avoids the disjunction problem. On Dretske’s view, a representational state must not only be causally dependent on the state of affairs it represents, but must serve a function for its containing system in virtue of this causal dependency. This extra condition motivates constraints on representational content that eliminate problematic disjunctive contents. Dretske’s theory is subject to some subtle criticisms that I will discuss in Section 6, but the Organism-Receptor account will preserve some of the teleological character of Dretske’s theory.

Ramsey’s most compelling objection to the receptor account, including Dretske’s sophisticated version, is that it justifies ascribing representational contents to states that are not, in fact, representational: smoke “represents” fire since the latter causes the former. Likewise, the firing pin of a gun “represents” whether the trigger is depressed, and rusting iron “represents” the presence of water and oxygen (138–47). Ramsey claims, plausibly, that these are absurd consequences. I find Ramsey’s *reductio*



Ambiguous figures. Left: The Necker cube. Right: The duck-rabbit (image from Jastrow 1899).

compelling, but reject a different premise than he does. Rather than conclude that cognitive scientists have a bad conceptual practice, I question whether his characterization of the receptor notion is a charitable understanding of what happens in cognitive science. After all, cognitive scientists do not generally claim that GOAT denotes *goats-or-sheep* (at least for competent judges of goathood), or that firing pins represent anything.

3. A Construal-based Notion of an Organism. I argue that something like the receptor notion can be salvaged if being a receptor is contextualized in terms of construal. Construal (also called “seeing-as”) is a judgment-like attitude whose semantic value can vary licitly independently of the state of affairs it describes. For example, we can construe an ambiguous figure like the Necker cube as if it were viewed from above or below, or the duck-rabbit as if it were an image of a duck or of a rabbit (Roberts 1988; see also Wittgenstein 1953). We can construe an action like

skydiving as brave or foolhardy, depending on which features of skydiving we attend to.

On a construal-based account of conceptual norms, a concept (e.g. REPRESENTATION) is ascribed relative to a construal of a situation. For example, perhaps I fear something only if I construe it as dangerous to me or detrimental to my ends (Roberts 1988). Daniel Dennett's (1987) intentional stance is a more familiar example: according to Dennett, a system has mental states if and only if we construe it in such a way that its behavior is explainable in terms of a belief-desire schema.

I propose that construing something as an organism involves construing it such that it has goals and behavior, and believing that it has mechanisms that promote those goals by producing that behavior. More precisely:

Organism-Construal. A subject a construes a system x as an organism in a context¹ c if and only if, in c ,

- (O1) a attributes a set of goals G to x ,
- (O2) a attributes a set of behaviors B to x ,
- (O3) a believes that the elements of B function to promote elements of G ,
- (O4) a believes that x possesses a set of mechanisms M , and
- (O5) a believes that the elements of M collectively produce the elements of B .

My main argument does not rely on all the details of Organism-Construal; it could be replaced by a different explication of what it is to see something as an organism. But Organism-Construal captures an intuitive notion of a critter. First of all, we normally take living critters to have goals, such as survival and reproduction, and behaviors that

¹ The relevant notion of a context is something like MacFarlane's (2014) "context of assessment."

promote those goals. However, Organism-Construal does not require that an organism really have goals (whatever that involves) or exhibit behavior (however that's distinguished from other performances). To see something as an organism according to Organism-Construal, the construing subject need only *attribute* goals to the system, and see some of its performances as behaviors that promote those goals. Such goals could include relatively specific aims such as locating food, getting out of the rain, or driving home. We sometimes also attribute goals and behaviors to non-living things, such as automated machines. For example, we might say that a robot vacuum has the goal of cleaning the floor, which it accomplishes by sucking up dust. Or I might say that my GPS navigation computer is trying to kill me, which it accomplishes by consistently giving me directions that lead me through strange, dangerous backroads. Condition (O3) is expressed in terms of belief instead of attribution, meaning that the construing subject must sincerely believe that an organism's putative behaviors function to promote its putative goals. When and insofar as someone construes a system in this way, the conditions (O1)–(O3) above are satisfied.

Conditions (O4)–(O5) require that the system's behavior be explainable by appeal to mechanisms. "Mechanisms" here should be understood in roughly the sense meant by the new mechanists (Machamer, Darden, and Craver 2000; Bechtel and Abrahamsen 2005; Craver 2007): organized structures of component parts and operations that produce a phenomenon, and the description of which is an explanatory aim of some scientific projects. Much explanation in biology and neuroscience plausibly follows a mechanistic model, and likewise in cognitive science. Daniel Weiskopf (2011) has argued that cognitive explanations are not properly mechanistic, but even on his view cognitive explanations are extremely similar to mechanistic ones, distinguishable only because the relationship between components of cognitive models and their physiological realizers is relatively opaque. Regardless, cognitive scientists use the word "mechanism" to refer to the referents of their models,

just as biologists and neuroscientists do. I am more moved by the similarities between the biological and the cognitive sciences than the differences. Therefore, like Catherine Stinson (2016), I acknowledge Weiskopf's concerns but nevertheless adopt the language of "mechanisms."

Not all of a system's mechanisms function to produce behavior. For example, biological organisms have metabolic and other mechanisms that maintain bodily integrity. Such mechanisms may need to function correctly as a background condition for the organism to behave, but scientists do not typically take behavioral patterns to be the explanandum phenomena of such mechanisms. Let us call mechanisms that do contribute to the explanation of behavior *behavioral mechanisms*. As for what it means for a system to "possess" a mechanism, a mereological criterion will do for now: the mechanism must be a part of the system. Condition (O5) is meant to limit the mechanisms in the set M to behavioral mechanisms.

So far so abstract; let's consider an example. The robot Herbert was designed to wander autonomously through the MIT robotics lab, avoiding obstacles, and collecting soda cans with its arm (Brooks, Connell, and Ning 1988). Herbert can be construed as an organism, even though it is not alive, as long as one (O1) attributes goals, like avoiding collisions and collecting soda cans, to Herbert, (O2) sees some of Herbert's performances as behaviors, (O3) believes that Herbert's behaviors promote its goals, and (O4) believes that Herbert possesses mechanisms that (O5) explain its behavior. Herbert does possess mechanisms for accomplishing goals; it is equipped with sensors, computers, and motors that coordinate its locomotion and its grasping arm. And most people readily anthropomorphize Herbert enough to see it as a goal-directed, behaving system (pace Adams and Garrison [2013], who insist that Herbert has its designers' goals, but no goals of its own). Anyone willing to engage in the imaginative attribution of goals and behavior to Herbert can see Herbert as an organism, even if on reflection they believe Herbert is not literally an organism. The

willingness to ascribe representations to a system plausibly waxes and wanes along with one's willingness to construe the system as an organism in something like the sense described above. There are psychological limits on the willingness to attribute goals and behaviors to systems relatively unlike animals, and these limits may vary between individuals.

4. The Receptor Notion Re-construed. Returning now to the receptor notion of natural representation, I suggest that it can be augmented in the following way:

Organism-Receptor. A state s represents a state of affairs p if

(R1) s is regularly and reliably caused by p , and

(R2) s is a functional state of a behavioral mechanism possessed by an organism.

Organism-Receptor is not a construal-based explication, but it depends on a construal-based account of ORGANISM. It preserves the spirit of Ramsey's receptor notion, with the added condition that representations be ascribed to parts of systems construed as organisms. Representation-ascriptions guided by Organism-Receptor inherit their plausibility from the plausibility of the corresponding construal of some system as an organism. Most accounts of cognitive representation require there to be a representational subject of some kind (e.g. Adams and Aizawa 2001; Rupert 2009; Rowlands 2010), and on Organism-Receptor the organism serves this role. We can constrain the acceptable contents of these representations by requiring they correspond to descriptions of p according to which p is relevant to the pursuit of an organism's goals. This appeal to goals is not ad hoc, since according to Organism-Receptor representations are ascribed to organisms, i.e. systems to which we've already attributed a set of goals. Thus, like Dretske's (1988) and Millikan's (1984)

teleofunctional accounts, this construal-based account addresses the disjunction problem by appealing to goals of organisms.

The metadiscursive job-description challenge is to provide criteria of ascription for representations, in virtue of which representation-ascriptions achieve some explanatory purpose. I have provided criteria of ascription, so what is their purpose? On Donald Davidson's (1963, 5) account of intentional action, actions are performed under the guise of a privileged description (or set of descriptions). Davidson flips the light switch in order to turn on the light, but not in order to alert the prowler outside (whose presence is unknown to Davidson) that he is home, though he also does the latter. Davidson calls this feature of action its "quasi-intensional character." Behavioral mechanisms also have something like a quasi-intensional character, since there are privileged descriptions that make explicit how they and their components contribute to an organism's capacity to pursue its goals. For example, edge-detecting cells in V1 fire in order to identify boundaries in an organism's environment, not to consume glucose, though they also do the latter. The use of representation-talk by cognitive scientists, as licensed by Organism-Receptor, is a way to habitually mark these privileged descriptions and distinguish them from other descriptions of the same states or events. And since cognitive science is concerned with the functional structure of behavior-coordinating mechanisms rather than other features of cognitive systems, it is easy to see why representation—even in this relatively thin sense—has always been the dominant theoretical perspective in cognitive science. This focus on quasi-intensional characterization may even be what makes the cognitive scientific perspective distinctive (on scientific perspectives, see e.g. Giere 2006).

The Organism-Receptor account provides us with resources to salvage the receptor notion from Ramsey's *reductio*. It is plausible to suppose that cognitive scientists generally ascribe natural representations to systems against an imaginative

background like this. After all, most cognitive science concerns the mechanisms of living systems, especially animals (except in computer science and some computational modeling, where the object of attention is a formal object like a connectionist network that is presumed to be analogous in some way to such a mechanism). Such systems are easily construed as organisms in the sense of Organism-Construal. Non-living things and even non-animals are in general more difficult to construe as organisms in that sense, since they are often perceived to lack goals, the capacity to behave, or both.

5. The Organism-Receptor Notion in Context. Consider a strong case of representation, like fly-detecting cells in frog visual cortex. We construe frogs as systems that exhibit goal-directed behavior and believe they possess mechanisms that explain that behavior. Frog visual cortex contains mechanisms that (along with other mechanisms) explain behaviors like fly-catching. When we identify cells in frog visual cortex that fire in response to the visual presence of flies (or fly-like objects), we ascribe representational properties to those cells. The contents we ascribe to representations in frog visual cortex are constrained by the goals we attribute to frogs. *That a small insect is present* is a suitable content because flies can be consumed for energy; *that a wiggly BB is present* does not have this significance for frogs, although BBs may be indistinguishable from insects by the mechanisms in the frog's visual cortex. Nevertheless, the relationship between fly-presence and the frog's goals provide a ground for privileging non-disjunctive descriptions of representational content.

The Organism-Receptor account also explains why liminal cases of representation, like the case of Herbert, are liminal. We can say that Herbert represents such states of affairs as the presence of obstacles and soda cans, because states of Herbert's sensors are regularly and reliably caused by those states of affairs.

And we can ascribe contents to representations by drawing on descriptions of Herbert's environment that relate to the goals we ascribe to Herbert. However, our willingness to take these representations seriously as natural representations that bear content intrinsically covaries with our willingness to take Herbert seriously as an organism. We are not as comfortable attributing genuine goals and behaviors to Herbert as we are attributing goals and behaviors to frogs.²

Finally, absurd cases like the firing pin can be excluded (for the most part) since guns are not easily construed as "organisms." Firearms are difficult to anthropomorphize, since they do not exhibit autonomous behavioral dynamics and we don't normally see them as having goals of their own. It is not *impossible* to ascribe goals to weapons or other tools, but the ascription of folk-psychological properties to tools, like the folk ascription of a bloodthirsty disposition to a sword, generally depends on the way a tool influences its users' behavior. (I suspect this dependence might offer some novel explanations of why Clark and Chalmers' [1998] extended cognition hypothesis is attractive to some.) The attribution of autonomous behaviors to tools like swords is fanciful. Perhaps we might imagine a tool exhibits psychic "behavior," but anyway we do not believe that swords possess mechanisms that produce this "behavior" (though if we did, such a construal would be more compelling). If the firing pin of a gun is not a component of a behavioral mechanism, it cannot represent anything according to the Organism-Receptor account.

So the Organism-Receptor account licenses an ascriptive practice that resembles the crude receptor notion when the role of construals is not made explicit. It is unusual in that it inverts Ramsey's preferred order of ascription: Ramsey wishes to

² Notably, Rodney Brooks himself does not claim that it is proper to ascribe representational capacities to Herbert (Brooks, Connell, and Ning 1988; Brooks 1991), but Brooks plausibly had in mind a more demanding account of representation.

ascribe cognitive structure to systems in virtue of their representational structure (see e.g. Ramsey, 222–235), whereas I suggest that we in fact ascribe representational structure in virtue of seeing a system as a system with goal-directed behavior, i.e. as a potentially cognitive system.

6. Worries. Since the Organism-Receptor account shares a certain teleological character with Dretske's account, I will discuss Ramsey's two most developed objections to Dretske, along with other worries specific to the Organism-Receptor account. First, Ramsey objects that Dretske's account is question-begging with regard to the job-description challenge. Roughly, teleological normativity (i.e. functioning and malfunctioning) is not sufficient to explain intentional normativity (i.e. representation and misrepresentation), and since Dretske provides no satisfying criteria for what it is for a state to function as a representation, he cannot bridge that gap (Ramsey 2007, 131–2). But the Organism-Receptor account has more resources than Dretske's teleofunctionalism. Construing a system as an organism involves construing it as exhibiting behavior, which allows us to distinguish behavioral mechanisms from other mechanisms. On the Organism-Receptor account, misrepresentations are malfunctions of behavioral mechanisms (like frog vision), but not of other mechanisms (like a frog's circulatory system or a gun's firing mechanism).

My reply invites a rejoinder: on the Organism-Receptor account the functional roles of representations will be extremely diverse, and representations will be common. They will not just include IO-representation and S-representation (roughly, information-processing relata and models for surrogative reasoning; Ramsey 2007, 68ff.), which Ramsey and most cognitive scientists regard as genuinely representational. They will also include more controversial varieties of "representation," such as Millikan's (1995) "pushmi-pullyu" representations: Janus-faced mechanistic components that simultaneously indicate a state of affairs and cause

an adaptive or designed response. In other words, representations will include what Ramsey calls “causal relays” like the firing pin in a gun, the inclusion of which in the extension of REPRESENTATION was the ground for his reductio! However, the absurd cases can be avoided. The firing pin case is excluded because guns are poor examples of organisms. And pushmi-pullyu representations include cases with significant intuitive appeal to many scientists, like the predator calls of vervet monkeys (Millikan 1995; cf. Seyfarth, Cheney, and Marler 1980). While this conception of representation has a more liberal extension than Ramsey is comfortable with, it is liberal enough to explain common representation-ascriptions in cognitive science without being so liberal as to countenance absurd cases like Ramsey’s firing pin, so I submit it is adequate to scientific practice.

Ramsey’s second objection is that Dretske is committed to a false principle: that if a component is incorporated into a mechanism because it carries information, then its function is to carry information (132–9). However, the Organism-Receptor account constrains the causal dependence criterion (R₁) by relying on construals of systems as organisms instead of teleofunctional commitments. The account I describe is not committed to Dretske’s principle, and therefore is not subject to this objection.³

Nevertheless, one might worry whether the organism criterion (R₂) is a suitable condition on representation-ascription. I suggested five conditions (O₁)–(O₅) on what can be seen as an organism, but conditions (O₁) and (O₂) are fairly unconstrained. There are psychological limitations on when goals or behaviors can be plausibly attributed to a system, but what are those limits? And what factors influence interpersonal variability in willingness to make these attributions? The reason this practice isn’t bonkers is that it coheres with the explanatory purpose of

³ Ramsey’s discussion is rich and worthy of deeper engagement than this, but for reasons of space I leave the matter here.

representation-ascriptions: to make explicit the quasi-intentional character of behavioral mechanisms. Nevertheless, we should hope that these psychological limitations are vindicated by more principled considerations. Criticism is warranted if scientists attribute goals and behaviors when they should not. There is some extant work on the proper norms ascribing goals to organisms (e.g. Shea 2013; Piccinini 2015, chap. 6), but little serious work on how to understand the concept of BEHAVIOR in the context of cognitive science. We should worry about the practice of ascribing natural representations if scientists construe things that are not cognitive systems as “organisms.” Indeed, we might indeed worry that many cognitive scientists misuse the concept COGNITION, given the intense disagreements over its extension (see e.g. Akagi 2017). However, my present aim is not to evaluate scientific practice, but to describe it faithfully (with the hope that a more satisfactory evaluation will follow).

Another worry about construal-based accounts is that they entail an unattractive anti-realism: if representations and their contents only exist relative to construals, they are mind-dependent rather than objective, right? This worry is unfounded. I am undertaking a modified version of Ramsey’s job description challenge: my aim is to describe the ascription of representations in virtue of which they serve an explanatory purpose, not to distinguish genuinely representational states from non-representational states. The Organism-Receptor account does not entail that representations exist relative to construals, only that they are *ascribed* relative to construals. My account is consistent with the existence of a first-order account of the metaphysics of representation that justifies this practice (or doesn’t). After all, the duck-rabbit can be construed as a duck even if it is not a duck, and nothing about that fact entails that ducks (or unambiguous images of ducks) are not real. The Organism-Receptor account describes a norm that plausibly guides human scientists with imperfect capacities for knowledge. But while my solution to the metadiscursive job description challenge is not inconsistent with Ramsey’s solution to

the first-order job description challenge, it is inconsistent with Ramsey's characterization of scientific norms for ascribing natural representations.

7. Conclusion. I began by observing the common worry that scientists ascribe representations more liberally than many philosophers are comfortable with, and in particular that scientists rely on an unsatisfactory "receptor" criterion. I sketched an account on which scientists ascribe natural representations only to components of mechanisms of systems construed as "organisms." Since in practice cognitive scientists attend almost exclusively to systems that are easily so construed, their behavior may appear to be guided by the crude receptor criterion whereas in fact it is guided by the Organism-Receptor criterion. However, while the Organism-Receptor account is still relatively liberal, a crucial difference between the two accounts is that the crude criterion has absurd consequences, whereas such consequences are eliminated or marginalized on the Organism-Receptor criterion. Since scientists do not in fact endorse these absurd consequences, I argue that the augmented criterion is a better hypothesis regarding norms for representation-ascription in cognitive science.

This proposal is not a comprehensive, new theory of representation, but it accomplishes two things. First, it provides argumentative resources for resisting the common worry that cognitive scientists use hopelessly liberal criteria for ascribing representations. Second, it offers a novel picture of practices for representation-ascription in the biological and behavioral sciences, one that is less pessimistic than Ramsey regarding conceptual rigor in cognitive science. The picture is not beyond criticism—in particular, it wants for a more detailed account of the grounds that warrant attributing behaviors and goals to systems. But since it is more faithful to our practice than Ramsey's it is likely to yield more productive suggestions for how to guide that practice into the future. I suggest that we safeguard conceptual rigor in cognitive science not by cleaving more faithfully to the representationalism of the

cognitive revolution, but by embracing role of construal in scientific inquiry, making it explicit, and subjecting it to reasoned criticism.

REFERENCES

- Adams, Fred, and Ken Aizawa. 2001. "The Bounds of Cognition." *Philosophical Psychology* 14:43–64.
- Adams, Fred, and Rebecca Garrison. 2013. "The Mark of the Cognitive." *Minds and Machines* 23:339–52.
- Akagi, Mikio. 2017. "Rethinking the Problem of Cognition." *Synthese*.
doi: 10.1007/s11229-017-1383-2.
- Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:421–41.
- Brooks, Rodney. 1991. "Intelligence without Representation." *Artificial Intelligence* 47:139–59.
- Brooks, Rodney, Jonathan Connell, and Peter Ning. 1988. "Herbert: A Second Generation Mobile Robot." *A.I. Memos* 1016:0–10.
- Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58:7–19.
- Craver, Carl F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- Cummins, Robert, and Pierre Poirier. 2004. "Representation and Indication." In *Representation in Mind: New Approaches to Mental Representation*, Edited by Hugh Clapin, Phillip Staines and Peter Slezak, 21–40. Amsterdam: Elsevier.
- Davidson, Donald. 1963. "Actions, Reasons, and Causes." *The Journal of Philosophy* 60:685–700.
- Dennett, Daniel C. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.

- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- . 1988. *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT/Bradford.
- Giere, Ronald N. 2006. *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Hubel, David H., and Torsten N. Wiesel. 1962. "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex." *The Journal of Physiology* 160:106–54.
- Jastrow, Joseph. 1899. "The Mind's Eye." *Popular Science Monthly* 54:299–312.
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Clarendon.
- Machamer, Peter K., Lindley Darden, and Carl F. Craver. 2000. "Thinking About Mechanisms." *Philosophy of Science* 67:1–25.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- . 1995. "Pushmi-Pullyu Representations." *Philosophical Perspectives* 9:185–200.
- Piccinini, Gualtiero. 2015. *Physical Computation: A Mechanist Account*. Oxford: Oxford University Press.
- Ramsey, William M. 2007. *Representation Reconsidered*. Cambridge: Cambridge University Press.
- Roberts, Robert C. 1988. "What Emotion Is: A Sketch." *Philosophical Review* 97:183–209.
- Rowlands, Mark. 2010. *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, MA: MIT Press.

- Rupert, Robert. 2009. *Cognitive Systems and the Extended Mind*. Oxford: Oxford University Press.
- Seyfarth, Robert M., Dorothy L. Cheney, and Peter Marler. 1980. "Monkey Responses to Three Different Alarm Calls: Evidence of Predator Classification and Semantic Communication." *Science* 210:801–3.
- Shea, Nicholas. 2013. "Naturalising Representational Content." *Philosophy Compass* 8:496–509.
- Stinson, Catherine. 2016. "Mechanisms in Psychology: Ripping Nature at Its Seams." *Synthese* 193:1585–614.
- Weiskopf, Daniel A. 2011. "Models and Mechanisms in Psychological Explanation." *Synthese* 183:313–38.
- Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. 3rd Ed. Trans. G.E.M. Anscombe. Eds. G.E.M. Anscombe and Rush Rhees. Oxford: Blackwell, 2001.