Going against the Grain: Functionalism and Generalization in Cognitive Science

MIKIO AKAGI

Abstract: Functionalism is widely regarded as the central doctrine in the philosophy of cognitive science, and is invoked by philosophers of cognitive science to settle disputes over methodology and other puzzles. I describe a recent dispute over extended cognition in which many commentators appeal to functionalism. I then raise an objection to functionalism as it figures in this dispute, targeting the assumption that generality and abstraction are tightly correlated. Finally, I argue that the new mechanist framework offers more realistic resources for understanding cognitive science, and hence is a better source of appeal for resolving disagreement in philosophy of science.

Word count: 4,985 words including abstract, headings, footnotes, and references.

1 Introduction. Functionalism is the doctrine that mental or cognitive states are functional states, whose identity conditions are articulable in terms of their characteristic inputs, outputs, and relations to other intermediate states. Functionalism was established as a central doctrine in the philosophy of cognitive science in the 1960s (Putnam 1967a, b, Fodor 1968), and though it has become less central to much contemporary discussion (Chemero & Silberstein 2008) it retains the notoriety of an orthodoxy in philosophy of mind (Buechner 2011) and in contemporary philosophy of cognitive science (Eliasmith 2002, Clark 2008, Sprevak 2009, Chalmers 2011). This remains true even though functionalism has been an embattled doctrine for decades (Block & Fodor 1972, Block 1980, Shagrir 2005, Godfrey-Smith 2008), has proliferated versions and variations (Levin 2013, Maley & Piccinini MS), and even though the canonical argument for functionalism—the argument from multiple realizability—has been subjected to a variety of criticisms (Batitsky 1998, Bechtel & Mundale 1999). This is all, importantly, to say nothing of other views that happen, unhappily, to be called "functionalism" in biology or in pre-behaviorist psychology (Cummins 1975, Sober 1985, Chemero 2009) but which have different intellectual lineages. My discussion concerns only Putnam's machine functionalism and derivative views. The persistence of functionalism is hardly a special case. It is the fate of many "received views," such as the belief-desire model of intentional action or the deductive-nomological model of explanation, to remain central to a literature despite decades of convincing criticism so long as there is no sufficiently dominant successor. The new mechanist view of explanation (Machamer, Darden, & Craver 2000, Bechtel & Abrahamsen 2005, Craver 2007) has recently achieved this status in the philosophies of the biological sciences, supplanting the deductive-nomological model and other law-subsumption models as a received view of explanation in those sciences. This is not to say that new mechanism is correct or uncontroversial, only that it has replaced other models of explanation as the primary target of interpretation and criticism.

Bearing in mind the tenacity of received views, my aim in this paper is not to simply poke more holes in the sinking ship of functionalism. Rather, I aim to promote an alternative vessel for philosophers of mind and cognitive science to pilot through choppy waters. To this end I will raise a "grain" objection to functionalism, based on the relationship between generalization and "fineness-of-grain." This objection is not a knock-down argument against all varieties or uses of functionalism. However, it needn't be, since functionalism is a sinking ship, and since my objection does apply to versions of functionalism that figure in notorious, recent disputes in the philosophy of cognitive science. I shall take as my example the controversy over extended cognition, especially its recent high-profile epicycle concerning the relation between extended cognition and functionalism (Rupert 2004, Clark 2008, Sprevak 2009). In the last section I will argue that new mechanism provides better resources for understanding variation between models in cognitive science, and for understanding the practice of generalization. In particular, mechanism is not vulnerable to the grain objection. I do not claim, of course, that mechanism is free from criticism or worries or that I have made clear what was once obscure. My aim, rather, is to motivate a change of focus in discussions of cognitive science from functionalism to mechanism.

2 Functionalism and Extended Cognition. Andy Clark and David Chalmers (1998) notoriously claim that cognition (like meaning) ain't all in the head. They argue that in certain cases the use of external props in some activities—a computer processor while playing some video games, one's notebook in carrying out one's plans for the day, perhaps one's partner in remembering past events—is such that those props should be considered parts of one's own cognitive economy, similarly to parts of one's brain. This claim has become known as the hypothesis of extended cognition (HEC). The most famous example concerns Otto, an older gentleman with a bad memory who uses a notebook to help him remember facts and plans.[1] In their argument, Clark and Chalmers appeal to what has become known as the "parity principle," which states that

> If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of a cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (Clark & Chalmers 1998, 8)

[1] The Otto example is originally an illustration not of extended cognition but of the extended mind, which is a distinct claim. Although this distinction is essential for charitably evaluating Clark and Chalmers' arguments, it is almost always ignored in the critical literature (even by Clark). Since I am not evaluating HEC here, I will ignore the distinction for ease of exposition.

One way to interpret this principle is as a corollary of functionalism: cognitive states are individuated by their functional relations (to inputs, outputs, and each other), and it is immaterial whether their realizers are located inside the brain or outside the body.[2] Thus, activities should count as cognitive processes if those body-external processes exhibit the same functional relationships (to inputs, outputs, and cognitive states) as other processes that we already happily consider cognitive processes. Of course, understood this way the parity principle only justifies a commitment to extended cognition if the functional relations are specified so that body-external activities and props do satisfy those specifications, and many cognitive and psychological processes can be specified in a variety of ways. Robert Rupert argues that Otto's notebook in the famous example cannot serve as a memory in part because it fails to satisfy the most fruitful functional description of human memory. Cognitive psychologists have documented many features of human memory—for example susceptibility to interference effects, generation effects, and conformity to the Rescorla-Wagner law (see Rupert 2004, 413–419). Since Otto's external "memory" does not exhibit these effects the parity principle does not license the attribution to Otto of extended cognitive processes (Rupert 2004). Fred Adams and Ken Aizawa (2001) argue for the same conclusion because Otto's use of his notebook must be described via *inter alia* relations to perceptual and motor intermediaries (he flips through the notebook, reads it, &c.), whereas canonical examples of internal memory are not related to perceptual and motor activities in the same way.

Mark Sprevak calls these objections the RAA (for Rupert, Adams, and Aizawa) objections. Sprevak suggests that "All varieties of functionalism contain a parameter that controls how finely or coarsely functional roles should be specified (how much should be abstracted and ignored)" (Sprevak 2009, 510). He observes that RAA trade on *fine-grained* differences between Otto's use of his notebook and canonical examples of memory. A coarse-grained functional description of memory might simply describe the relations between past perceptions and actions and future behaviour, but not describe memory as e.g. exhibiting interference or generation effects, or as obeying the Rescorla-Wagner law. Fine-grained functional descriptions may specify these relations, but are objectionable because they conflict with the common intuition that there could be Martians who have cognitive processes but whose cognitive architecture is distinctly different from ours. Such Martians, unlike us, may not exhibit interference or generation effects, and may even store information by manipulating ink-marks on paper inside of their brains, and retrieve it by reading the marks back with photosensitive organs. The "Martian intuition" is that while this is an alien form of memory, it is memory nonetheless. Since such Martians have memory, and their memory may have the same fine-grained functional description as Otto's use of his notebook, the parity principle demands that we consider Otto's a case of extended cognition. Thus Sprevak argues that functionalism implies HEC.

Unfortunately for the defenders of HEC, however, Sprevak argues that coarse-grained functional descriptions are no more acceptable, for the parity principle is less restrictive than Clark and Chalmers anticipate. Since we can imagine far-fetched Martian minds, the parity principle

[2] This is accepted by most of Clark and Chalmers' critics (Adams & Aizawa 2001, Rupert 2004, Sprevak 2009), and is almost certainly not the best interpretation of the parity principle.

licenses a radical form of HEC. For example, we might imagine Martian minds that are embedded with factual information that must be retrieved with effort, so that this process has functional parity with the activity of looking up information in a library. Such possibilities seem to license radical cases of extended cognition: that contents of volumes in a library are beliefs of any person in the library, or that being in possession of a graphing calculator gives one a knowledge of integral calculus (517–518). These consequences, Sprevak argues, are absurd, and justify a *reductio* of radical HEC and, since it entails radical HEC, of functionalism.

3 Going against the Grain. The dispute over HEC is not seen by its partisans as an idle philosophical discussion, but as a battle for the soul of cognitive science. If HEC is true, it is claimed, it has dramatic consequences for the way cognitive scientists conduct their research. Hence, both defenses and criticisms of HEC draw on empirical results and claims about theory-choice in science (e.g. Clark & Chalmers 1998, Rupert 2004, Adams & Aizawa 2008, Clark 2008, Rowlands 2010). The fact that so many of the arguments concerning HEC trade on interpretations of functionalism reveals the belief of many that functionalism provides a suitable framework for understanding cognitive scientific models. Disagreements about HEC force a discussion of what precisely the laws of cognitive science are—both what their proprieties are with respect to generalization, and what phenomena should be investigated and accommodated in order to construct those laws. In Putnam-style functionalism, functional descriptions (i.e. via Ramsey sentences, cf. Lewis 1972) operate as laws characterizing mental states. Putnam's proposal aspires explicitly to generalization over diverse kinds of system—e.g. about pain in primates and also in cephalopods. The arguments that arise in connection with the RAA objections to HEC concern whether descriptions of e.g. memory generalize over head-internal vehicles and extended vehicles (like Otto's notebook-use).

Like Clark and Sprevak, Hilary Putnam is wary of psychological chauvinism (human-specificity). A type-physicalist account of pain (Place 1956, Smart 1959), like the simplistic conjecture that pain is the activation of C-fibers, denies without motivation that animals that lack C-fibers have pain-states. In Putnam's canonical argument for functionalism, functionalism achieves *generality* by proposing *abstracted* descriptions that omit physiological and other details. Sprevak's grain parameter makes this maneuver more explicit by proposing a continuum of descriptions that are increasingly abstract, in the sense of omitting detail, and therefore increasingly general. (Some may object to the use of the word "abstract" as the complement of "detailed," rather than of "concrete." However, Sprevak uses the word this way and there is ample precedent for his doing so (Levy & Bechtel 2013).) Consider a toy functional description of pain: pain is caused by tissue damage, and causes stress, increased metabolic activity, and evasion of the damaging stimulus. This description denotes processes in a variety of complex organisms, including cephalopods (which lack C-fibers). Elements can be added to this description to make it more fine-grained, and to denote processes in progressively more restricted classes of organisms. For example, if pain also tends to cause excited vocalization, then creatures like cephalopods which do not vocalize will not satisfy this more fine-grained description of pain.

However, Sprevak's grain parameter is not an effective way of capturing variation between cognitive models. In the space of models that cognitive scientists actually produce, generality-specificity and abstraction-detail are independent dimensions of variation. By way of example, I shall mention two cognitive models in which generality and abstraction are dissociated. The first, the motor theory of speech perception (Liberman et al. 1967, Liberman & Mattingly 1985), is quite abstract but highly specific to humans. The motor theory claims that "perceiving speech is perceiving gestures," and more specifically that the recognition of phonemes and words in natural language is mediated by processing in the motor system, namely motor processing that also governs the articulation of speech in the vocal tract. There are animals other than humans that can identify phonemes and words—dogs commonly learn to recognize some words, and chinchillas have been trained to distinguish natural language phonemes (Kuhl & Miller 1978)—but since they do not have the relevant vocal capacities they most likely exhibit this capacity exclusively by recognizing auditory patterns, whereas humans do not. Nevertheless, the motor theory is quite abstract—it specifies that speech perception depends on structures that govern vocalization. While there are more detailed claims about how this dependency manifests in humans (e.g. McGurk & MacDonald 1976), all of them are consistent with the motor theory.

On the other hand, feature-detector models of vision (e.g. Barlow 1953, Hubel & Wiesel 1962) are detailed, but general. Even normalization-based models of feature-detection in particular (Heeger 1992, Carandini & Heeger 1994), which are described by Mazviita Chirimuuta (2014), are quite general. These models describe sensitivity to contrasts, edges, &c. in early stages of visual processing, and unify evidence about the receptive fields of individual neurons as well as computational models of their response dynamics. On Heeger's normalization model, neurons in visual cortex respond linearly to excitatory input from the lateral geniculate nucleus, but inhibit each other "laterally" according to an equation. The terms of the equation stand for properties and activities of individual neurons and populations of neurons. This model can be integrated into conjectures about the gross architecture of visual cognition (Marr 1982), and features in the "standard model" of primary visual cortex (Rust & Movshon 2005). However, even without supplementation with other models of visual processing the normalization model makes quantitative predictions about neuronal activity and has a well-specified physiological interpretation. Nevertheless, despite the level of detail in contemporary feature-detector models, they do not apply only to humans. Early evidence for normalization was gathered largely from cats and frogs, and the models may generalize to all vertebrate vision.

I am sympathetic to Sprevak's conclusion that functionalism is false, however functionalism is in worse shape than he acknowledges. His argument presupposes that we can manipulate abstraction from detail like the mesh of a sieve to sift the chauvinistic cognitive models from the liberal models. However, the motor theory of speech perception and feature-detector models of visual processing illustrate the double dissociation between abstraction from detail and generality over diverse kinds of cognitive systems. If the grain parameter is supposed to track degrees of abstraction from detail, then it fails to simultaneously track generalizability in cognitive models. If it is meant to track both, it fails to accurately capture the variation in cognitive models. Either way it incorporates false presuppositions about the character of the variety in cognitive models.

4 Generality without Laws. The problems with functionalism that are made explicit in the "grain" objection are inherited from the covering-law view of explanation and generalization that was popular throughout the twentieth century. On that conception, generalization is achieved by subsuming many phenomena under a common description (expressing a "covering law"). However, the covering-law view has in recent years been supplanted by the new mechanist view of explanation, at least in the biological sciences. The mechanists hold that many scientific explanations, including a preponderance of explanations in the biological sciences, are achieved by specifying models of mechanisms. The extension of the mechanist view to cognitive science requires the suppression of certain controversial assumptions developed for biological contexts (especially certain assumptions of Craver 2007, see Weiskopf 2011, Chirimuuta 2014), but not all mechanists make these assumptions (cf. Machamer, Darden, & Craver 2000, Bechtel 2008).

Let us suppose that the primary explananda of cognitive science are intelligent behaviors or cognitive capacities. Intelligent behavior is behavior that is sensitive to the circumstances of an organism and that can be rationalized by its relation to a goal of the organism; cognitive capacities are those that are exhibited in intelligent behavior. A cognitive mechanism, then, is a structure of component entities and component operations that are organized such that they produce intelligent behavior (adapted from Machamer, Darden, & Craver 2000, Bechtel & Abrahamsen 2005, Craver 2007). The entities that figure in cognitive mechanisms are things like representations, modules, brain areas, populations of neurons, or idealized "neurons" in artificial neural networks. Characteristic operations in cognitive mechanisms are processing operations on or between those entities: transformations of representations, computational interactions between modules and brain areas, activation and inhibition of neuron populations, and interactions between artificial neurons as specified by connection weights. The organization of these entities and operations into mechanisms is usually represented by graphs, but can be specified more or less completely by groups of equations or descriptions of relations between components. Cognitive models are models of how mechanisms produce cognitive capacities (possibly or actually), and functional roles can be assigned to components of the models according to how those components contribute to the mechanism's production of that capacity (roughly as described in Cummins 1975). Cummins-style functional roles, however, are not functional descriptions; they describe a component's contribution to a capacity rather than conferring identity conditions in virtue of relations to input, outputs and intermediate states, and are thus independent from Putnam-style functionalism (see Craver 2001 for a discussion of Cummins-functions and neo-mechanism).

If the mechanist framework is to overcome functionalism's difficulties with generalization, it must provide an alternative to the covering-law framework, or even a covering-model framework

(Bechtel & Abrahamsen 2005, Craver 2007, 66–70). After all, it is now widely believed that the biological and social sciences have no true laws. William Bechtel and Adele Abrahamsen (2005) suggest that mechanistic explanations are generalizable not because the target systems are *identical* in the relevant respects, but because they are *similar*:

> The need to invoke similarity relations to generalize mechanistic explanations seems to be a limitation of the mechanistic account. But in fact it may be the mechanistic account that provides a better characterization of how explanations are generalized in many sciences. Laws are generalized by being universally quantified and their domain of applicability is specified by the conditions in their antecedents. On this account, no instance better exemplifies the law than any other. But in actual investigations of mechanisms, scientists often focus on a specific exemplar when first developing their accounts. (Bechtel & Abrahamsen 2005, 438)

The claim that generalization is based on similarity to exemplars is less satisfying than the picture of subsumption under a covering law. Bechtel and Abrahamsen's claims do little to constrain the practice of licit generalization, and their observation that scientists "seem to have an intuitive sense" of how to generalize is distinctly unsatisfying (ibid). However, given the lack of universal or exceptionless laws in the biological sciences, a more complicated conception of generality is needed. The need to be more specific about "similarity"-based generalization is not a drawback of the mechanist framework, but a demand for further research by philosophers of science.

The mechanist framework offers richer resources than functionalism for constrained similarity-comparisons. First, mechanism models are more structured than functional descriptions. Functional descriptions might be structured according to independent predicates or conjuncts inside the scope of the quantifier in a Ramsey sentence. In comparing two mechanisms, one can appeal to similarities and differences between the sets of entities, of operations, their properties, or in their organization. Importantly, the result of such comparisons is not a judgment that mechanisms described by different models are simply the same or different, but that they are similar in certain respects and dissimilar in others. Frequently, a model may apply but with modifications, with the consequence that insights are gained both for the new and for the original target systems. For example, the two visual streams hypothesis (Milner & Goodale 2006) was developed for primate visual systems, primarily with data from humans and macaques, but comparisons of primates and other organisms such as frogs enrich the model (see e.g. Goodale & Humphrey 1998, 183–185) and provide a framework around which similar models can be developed for most vertebrates (Jeannerod & Jacob 2005, 301). Generalization here is achieved through comparisons to exemplars with acknowledgement of differences, not subsumption under a common description. This is mechanistic generalization by, if you like, functional similarity, but not functional identity in Putnam's sense. Since similarity-based generalization like this does not presuppose that generality and abstraction from detail are correlated. Thus the motor theory of speech perception and the normalization model of visual feature-detection are not anomalies in the mechanist framework.

It might be possible to provide similarity-based generalizations of functionally-individuated kinds, but such a strategy is not pursued by those who appeal to functionalism in order to settle other questions in philosophy of cognitive science. For example, the strategy is not pursued by Adams and Aizawa, Rupert, or Sprevak in their criticism of HEC, who instead seek categorical descriptions of mental or cognitive processes. Clark and Chalmers appeal to a relatively abstract specification of memory to argue that Otto's notebook functions as a part of his memory. RAA appeal to relatively detailed specifications of memory to argue that he does not. An ecumenically-minded theorist might suggest that alternative specifications—some detailed and some abstract—delineate various dimensions of similarity and difference between paradigmatic memory and Otto's notebook-augmented memory. However, such a proposal must specify how membership is decided for the set of admissible descriptions for a term. The main products of cognitive scientific research (apart from philosophical research) are models, not functional descriptions. The mechanist framework provides a more natural resource for appeal in philosophy of cognitive science than an unarticulated successor to functionalism. In general the place for functionalism as a resource for appeal in philosophy must be reevaluated.

5 Conclusion. My intention in this paper was to show that the assumptions of functionalism are inappropriate for thinking clearly about cognitive science. To this end I described some discussion of the RAA objections to HEC, and claimed that Sprevak's "grain parameter" makes explicit an assumption that features in the motivating arguments for functionalism: that abstraction from detail and generality are correlated features of cognitive models. This assumption is false, so functionalism is an inappropriate framework for characterizing cognitive models and for settling disputes about cognitive science, like the dispute over HEC, that turn on generality. I suggested, following a suggestion by Bechtel and Abrahamsen, that where the functionalist framework hides the complexity in cognitive scientists' practice of generalization, the mechanist literature provides a more fruitful framework for exploring that complexity. I have not argued that mechanism settles whether HEC is true or false. However, if disagreements about HEC are to be a battle for the soul of cognitive science, the proper battleground is over what kinds of mechanisms are cognitive ones, not over functionalist descriptions of mental states (cf. Walter 2010). The mechanist framework does not provide us with resources for determining the identity conditions of cognitive phenomena like belief and memory, as the functionalist framework does. However, cognitive scientists do not take conformity to their models as a criterion of exhibiting a phenomenon. For example, psychologists do not claim that exhibiting interference effects is a necessary condition on memory. That a system does not exhibit interference effects implies that memory models that do exhibit such effects must be modified in order to be generalized to that target system, not that the target system lacks genuine memory. It is therefore peculiarly contentious for philosophers to appeal to these models in order to settle the identity conditions for cognitive phenomena under the guise of being scientific. The contentious nature of this form of argument is no doubt obscured by the common belief that functionalism is an orthodoxy of cognitive science.

REFERENCES

Adams, Fred, and Ken Aizawa. 2001. "The Bounds of Cognition." *Philosophical Psychology* 14: 43–64.

---. 2008. *The Bounds of Cognition*. Malden, MA: Blackwell.

Barlow, Horace D. 1953. "Summation and Inhibition in the Frog's Retina." *The Journal of Physiology* 119: 69–88.

Batitsky, Vadim. 1998. "A Formal Rebuttal of the Central Argument for Functionalism." *Erkenntnis* 49: 201–220.

Bechtel, William. 2008. "Mechanisms in Cognitive Psychology: What Are the Operations?" *Philosophy of Science* 75: 983–994.

Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 421–441.

Bechtel, William, and Jennifer Mundale. 1999. "Multiple Realizability Revisited: Linking Cognitive and Neural States." *Philosophy of Science* 66: 175–207.

Block, Ned. 1980. "Troubles with Functionalism." In Ned Block (ed.), *Readings in Philosophy of Psychology*. Cambridge, MA: Harvard University Press, pp. 171–184.

Block, Ned, and Jerry A. Fodor. 1972. "What Psychological States Are Not." *Philosophical Review* 81: 159–181.

Buechner, Jeff. 2011. "Not Even Computing Machines Can Follow Rules: Kripke's Critique of Functionalism." In Alan Berger (ed.), *Saul Kripke*. Cambridge: Cambridge University Press, pp. 343–367.

Carandini, Matteo, and David J. Heeger. 1994. "Summation and Division by Neurons in Primate Visual Cortex." *Science* 264: 1333–1336.

Chalmers, David J. 2011. "A Computational Foundation for the Study of Cognition." *Journal of Cognitive Science* 12: 323–357.

Chemero, Anthony. 2009. *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.

Chemero, Anthony, and Michael Silberstein. 2008. "After the Philosophy of Mind: Replacing Scholasticism with Science." *Journal of Philosophy* 75: 1–27.

Chirimuuta, Mazviita. 2014. "Minimal Models and Canonical Neural Computations: The Distinctness of Computational Explanation in Neuroscience." *Synthese* 191: 127–153.

Clark, Andy. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.

Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58: 7–19.

Craver, Carl F. 2001. "Role Functions, Mechanisms, and Hierarchy." *Philosophy of Science* 68: 53–74.

---. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.

Cummins, Robert. 1975. "Functional Analysis." *The Journal of Philosophy* 72: 741–765.

Eliasmith, Chris. 2002. "The Myth of the Turing Machine: The Failure of Functionalism and Related Theses." *Journal of Experimental and Theoretical Artificial Intelligence* 14: 1–8.

Fodor, Jerry A. 1968. "The Appeal to Tacit Knowledge in Psychological Explanations." *The Journal of Philosophy* 65: 627–640.

Godfrey-Smith, Peter. 2008. "Triviality Arguments against Functionalism." *Philosophical Studies* 145: 273–295.

Goodale, Melvyn A., and Keith G. Humphrey. 1998. "The Objects of Action and Perception." *Cognition* 67: 181–207.

Heeger, David J. 1992. "Normalization of Cell Responses in the Cat Striate Cortex." *Visual Neuroscience* 9: 181–197.

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: 106–154.

Jeannerod, Marc, and Pierre Jacob. 2005. "Visual Cognition: A New Look at the Two-Visual Systems Model." *Neuropsychologia* 43: 301–312.

Kuhl, Patricia K., and James D. Miller. 1978. "Speech Perception by the Chinchilla: Identification Functions for Synthetic Vot Stimuli." *Journal of the Acustical Society of America* 63: 905–917.

Levin, Janet. 2013. "Functionalism." In *Stanford Encyclopedia of Philosophy*, ed Edward N. Zalta. http://plato.stanford.edu/archives/fall2013/entries/functionalism/.

Levy, Arnon, and William Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80: 241–261.

Lewis, David K. 1972. "Psychophysical and Theoretical Identifications." *Australasian Journal of Philosophy* 50: 249–258.

Liberman, Alvin M., Franklin S. Cooper, Donald P. Shankweiler, and Michael Studdert-Kennedy. 1967. "Perception of Speech Code." *Psychological Review* 74: 431–461.

Liberman, Alvin M., and Ignatius G. Mattingly. 1985. "The Motor Theory of Speech Perception Revised." *Cognition* 21: 1–36.

Machamer, Peter K., Lindley Darden, and Carl F. Craver. 2000. "Thinking About Mechanisms." *Philosophy of Science* 67: 1–25.

Maley, Corey, and Gualtiero Piccinini. MS. "Get the Latest Upgrade: Functionalism 6.3.1."

Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge, MA: MIT Press, 2010.

McGurk, Harry, and John MacDonald. 1976. "Hearing Lips and Seeing Voices." *Nature* 264: 756–748.

Milner, A. David, and Melvyn A. Goodale. 2006. *The Visual Brain in Action. 2nd Ed.* Oxford: Oxford University Press.

Place, U.T. 1956. "Is Consciousness a Brain Process?" *British Journal of Psychology* 47: 44–50.

Putnam, Hilary. 1967a. "The Mental Life of Some Machines." In Hector-Neri Castañeda (ed.), *Intentionality, Minds, and Perception*. Detroit: Wayne State University Press, pp. Reprinted in Putnam, 1975, *Mind, Language and Reality* (Cambridge: Cambridge University Press), pp. 408–428.

---. 1967b. "Psychological Predicates." In William H. Capitan and Daniel Davy Merrill (eds.), *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press, pp. 37–48. Reprinted in Putnam, 1975, *Mind, Language and Reality* (Cambridge: Cambridge University Press), pp. 429–440. As "The Nature of Mental States.".

Rowlands, Mark. 2010. *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, MA: MIT Press.

Rupert, Robert. 2004. "Challenges to the Hypothesis of Extended Cognition." *Journal of Philosophy* 51: 389–428.

Rust, Nicole C., and J. Anthony Movshon. 2005. "In Praise of Artifice." *Nature Neuroscience* 8: 1647–1650.

Shagrir, Oron. 2005. "The Rise and Fall of Computational Functionalism." In Yemima Ben-Menahem (ed.), *Contemporary Philosophy in Focus: Hilary Putnam*. Cambridge: Cambridge University Press, pp. 220–250.

Smart, J.J.C. 1959. "Sensations and Brain Processes." *Philosophical Review* 68: 141–156.

Sober, Elliott. 1985. "Panglossian Functionalism and the Philosophy of Mind." *Synthese* 64: 165–193.

Sprevak, Mark. 2009. "Extended Cognition and Functionalism." *The Journal of Philosophy* 106: 503–527.

Walter, Sven. 2010. "Cognitive Extension: The Parity Argument, Functionalism, and the Mark of the Cognitive." *Synthese* 177: 285–300.

Weiskopf, Daniel A. 2011. "Models and Mechanisms in Psychological Explanation." *Synthese* 183: 313–338.