Prosthetic Models

Carl F. Craver[†‡]

What are the relative epistemic merits of building prosthetic models versus building nonprosthetic models and simulations? I argue that prosthetic models provide a sufficient test of affordance validity, that is, of whether the target system affords mechanisms that can be commandeered by a prosthesis. In other respects, prosthetic models are epistemically on par with nonprosthetic models. I focus on prosthetics in neuroscience, but the results are general. The goal of understanding how brain mechanisms work under ecologically and physiologically relevant conditions is narrow compared to the search for maker's knowledge about how the brain can be made to work for us.

1. Introduction. Recent advances in building prosthetic sensory systems, brain-machine interfaces, and artificial cells have energized research programs to build prosthetic devices that can replace central brain regions. These research programs have possibly far-reaching medical and social implications. My focus is on their epistemic value. What, if anything, does the effort to build a prosthesis contribute to the search for neural mechanisms over and above what more familiar models and simulations contribute?

After defining my terms (sec. 2), I discuss the relative advantages of simulation models (sec. 3.1) and prostheses (sec. 3.2) for testing the models implemented in the simulations and prosthetic devices. I introduce the term "affordance validity" to describe the extent to which the target system has features that can be commandeered by a prosthesis to perform some function. I argue that the effort to build a prosthetic model allows a

†To contact the author, please write to: Department of Philosophy, Philosophy-Neuroscience-Psychology Program, Washington University in St. Louis, One Brookings Drive, Wilson Hall, St. Louis, MO 63130; e-mail: ccraver@artsci.wustl.edu.

‡Thanks to Anna Alexandrova, Masviita Chirimuuta, Dennis Des Chene, David Kaplan, Dan Moran, Wendy Parker, Susan Sterrett, and the members of the Washington University Philosophy of Science Co-op for comments and discussion.

Philosophy of Science, 77 (December 2010) pp. 840–851. 0031-8248/2010/7705-0025$10.00
Copyright 2010 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/656822 Published online by Cambridge University Press

decisive test of affordance validity but offers no distinct advantages for assessing the model's phenomenal and mechanistic validity. These conclusions follow from the fact that the engineering objectives of building a neural prosthesis and the epistemic objectives of building phenomenally and mechanistically valid models need not coincide (sec. 4).

2. Models, Simulations, and Prostheses. Models are descriptions. Targets are the things they describe. *Phenomenal models* of a mechanism describe the function relating a mechanism's inputs to its outputs (see Mauk 2000). *Mechanistic models* describe the parts, activities, and organization that explain why the input-output relationship holds (Dayan and Abbott 2001; Craver 2007). Hodgkin and Huxley's (1952) model of changes in membrane conductance during an action potential is purely phenomenal; it describes a curve fit to data produced by holding a neuron's membrane voltage fixed, measuring the resulting ionic currents, and inferring from the changes in current flow what the change of membrane conductance to the different ions must be. As Hodgkin and Huxley insist, their model does not describe the mechanisms by which voltage changes conductance (Bogen 2005; Craver 2006, 2007, 2008). More recent work on how voltage-sensitive ion channels in the membrane open and close as voltage changes offers mechanistic models that predict the curvilinear relations that Hodgkin and Huxley described (see Hille 1992).

A *simulation* is a process taken to mimic the relevant features of a target (cf. Hartman 1996; Guala 2002). As above, *phenomenal simulations* of a mechanism mimic the mechanism's input-output function, and *mechanistic simulations* mimic its parts, their activities, and their organization.[1] When a mechanistic model is written in code and implemented on a computer, the computer running the program simulates the target. Several simulations of the action potential are available online (see Benzania 1997–2003; Touretsky et al. 2008). The program behind the scenes is a model that represents the mechanism's components closely enough to mimic the response to interventions that change those components.

A *prosthesis* is a device designed to replace or restore the function of some biological component. Not all prostheses simulate the mechanisms they are designed to replace. Prosthetic legs for sprinters, for example, are made of L-shaped flexible materials that bear little or no resemblance to any biologically human leg. A *phenomenal prosthesis* might mimic the behavior of the mechanism using a different kind of mechanism entirely. A *mechanistic prosthetic model*, in contrast, is an engineered simulation

1. Following Korb and Mascaro (2009), SIM is a mechanistic simulation of a target, T, when SIM and T are both mechanisms, and a model representing the entities, activities, and organizational features of T also applies to SIM.

https://doi.org/10.1086/656822 Published online by Cambridge University Press

of a mechanism causally integrated into a biological system to replace the function of the target. Mechanisms are typically themselves components in higher-level systems. The target mechanism interacts with the other components in the system through *interfaces* (its inputs and outputs). The goal of building a prosthetic model is to replace the behavior of the missing part in context while preserving the behavior of the system as a whole. Prosthetic mechanistic models, in short, are simulations of a target mechanism that interface with a target system.

For example, computer simulations of mathematical models such as the Hodgkin-Huxley equation are now used to provide real-time voltage-dependent manipulations of neurons (Prinz, Abbott, and Marder 2004; Destexhe and Bal 2009). Some hybrid models (as these prosthetic models are called) mimic channel populations and synaptic inputs. Others mimic the behavior of entire cells linked via electrodes into circuits of neurons.

Five evaluative dimensions are especially relevant for assessing models, simulations, and prostheses: completeness, verification, phenomenal validity, mechanistic validity, and affordance validity. First, models vary in their *completeness*. All models and simulations of mechanisms omit details to emphasize certain key features of a target mechanism over others. Models are useful in part because they commit such sins of omission (Mauk 2000).

Second, simulations and prostheses can vary in the extent of their *verification*, that is, in the extent to which the simulation or prosthesis faithfully implements the intended model. Physical features of the computational device, for example, frequently impose constraints on what types of function can be implemented and on the degree of fidelity with which they can be implemented.

Models, simulations, and prostheses also vary in their *validity*, in the extent to which the model and the world match one another in relevant respects. Three types of validity are relevant. A model or simulation is *phenomenally valid* to the extent that its input-output function is relevantly similar to the input-output function of the target. A model or simulation is *mechanistically valid* to the extent that the parts, activities, and organizational features represented in the model are relevantly similar to the parts, activities, and organizational features in the target.

Finally, a model is *affordance valid* to the extent that the behavior of the simulation could replace the target in the context of a higher-level mechanism. The term "affordance" is chosen for its Gibsonian resonance: to address whether the higher-level mechanism "affords," or makes available, mechanisms with which the prosthesis can interface (Cummins 1975; Craver 2001). I show below that prosthetic models allow for decisive tests of one dimension of affordance validity and that prosthetic models set a

https://doi.org/10.1086/656822 Published online by Cambridge University Press

very high bar for the fidelity criteria for phenomenal match between the model and the target.

3. Prostheses for Testing Models of Target Mechanisms. My central question is what, if anything, building a prosthetic mechanistic model adds to our confidence that we have a valid mechanistic model over and above the degree of confidence provided by models and simulations alone.

Think of evidence as a finding that shapes (or constrains) the space of possible mechanisms for a given phenomenon. Points in the space are models of a mechanism. Regions represent families of similar models. Some evidence adds or removes points or regions of the space by suggesting new models and ruling out others. Other evidence redistributes probabilities over the space. Validity as defined above is a matter of fit (phenomenal, mechanistic, or affordance) between a model and the world.

3.1. Simulation and Validity. To answer our central question, we must first ask how simulation helps to constrain the space of possible mechanisms. First, simulations can be used to test a model's completeness. Models with gaps or poorly specified steps cannot be "run" because a program built to the model's specifications could not get past the black or gray boxes. The machine is missing crucial parts and functions. If one builds a simulation according to the blueprint provided by a model and the simulation does not work at all (it does not run), then as long as the simulation is an implementationally verified realization of the model, the model does not accurately describe how the target mechanism works. So, verified implementation of a model favors that model relative to models that cannot be implemented or whose implementability remains a matter of conjecture. Many box and arrow diagrams in cognitive science await computer simulations because they are too imprecisely specified or gappy to be implemented. This was Marr's (1981) central complaint about cognitive science.

Second, simulations can be used to test a model's phenomenal adequacy. The modeler compares the results of experiments performed on the simulation to the results of experiments performed on the target (see Parker 2009). Experiments on simulations involve fixing the values of some parameters in the model and checking the effects of that intervention on the values of other variables. The modeler then manipulates and measures the target mechanism to see if it behaves similarly when it works or to see if it responds similarly when the experimenter intervenes to set the target variables to particular values. Simulation allows one effortlessly to perform laborious calculations. It is therefore an efficient tool for exploring a model's empirical commitments. There is no principled difference, however, between simulating a model on a computer and simulating

https://doi.org/10.1086/656822 Published online by Cambridge University Press

it by putting pencil to paper. Huxley, for example, worked out the consequences of their model with an adding machine.

The general point is that, if one builds the simulation according to a blueprint provided by the model, the simulation is implementationally verified, and the simulation does not behave like the target, then the model is not phenomenally valid. It can be removed from the space of possible models.

A simulation might behave like the target under standard conditions or under all relevant conditions within the target system but fail to do so under nonstandard conditions. A model or simulation is *narrowly phenomenally valid* to the extent that its input-output relationship is relevantly similar to the target input-output relation under standard or "normal" input conditions. A model is *widely phenomenally valid* to the extent that it mimics the target's behavior when the inputs are outside of the standard range as well. Widely phenomenally valid simulations are required to test for mechanistic validity given that one expects the simulation of a mechanism to break and malfunction in ways that are relevantly similar to the conditions under which the target mechanism breaks and malfunctions. This difference is crucial for distinguishing how-possibly from how-actually models (Craver 2007, chap. 4). It is also crucial for understanding how prostheses and simulations differ.

Turn now to a simulation's mechanistic validity. The space of implementable and phenomenally adequate models for a given phenomenon is typically very large. Phenomena are *multiply realizable* in lower-level mechanisms. Multiple realizability obstructs the inference from a model's phenomenal validity to its mechanistic validity. The space of phenomenally adequate simulations might well be too large and heterogeneous to provide any assurance that the mechanistic features of a phenomenally adequate simulation are relevantly similar to the mechanistic features of the target. As the set of phenomenally adequate models grows, the confirmatory value of phenomenal validity with respect to mechanistic validity diminishes.

For this reason, many simulations in neuroscience are intended to be more or less mechanistically valid or plausible. Researchers often use simulations to test a model's mechanistic validity. They compare the changes of internal variables in the simulation to the changes observed among the measurable variables in the target mechanism. Again, simulation per se adds no epistemic force that models do not already have. Simulation simply takes the math out of our hands.

Consider now the use of simulations to test affordance validity. For simulations, affordance validity amounts to phenomenal validity. One forms a hypothesis about the relevant input-output function exhibited by the target mechanism and builds a simulation to mimic that input-output

https://doi.org/10.1086/656822 Published online by Cambridge University Press

function. As a result, simulations count as tests of affordance validity only on the assumption that one has correctly identified the relevant interfaces between the target mechanism and its containing system and that one's instruments reliably detect the relevant properties in the system.

Models and simulations almost always differ in some respects from their targets. The question of fidelity arises: How closely must they match to be judged phenomenally valid (see Weisberg 2007)? Prosthetic models set a high bar for fidelity criteria. The epistemic act of comparing the simulation's input-output relation to the target's input-output relation is replaced in prosthetic modeling by a real-time causal interface between the simulation and the target.

3.2. Building a Prosthesis and Knowing How It Works. So now, what constraints do prosthetic models add to the space of possible mechanisms beyond those provided by simulations?

Consider mechanistic validity first. Prosthetic models at their most biologically realistic are engineered simulations. As such, they inherit the epistemic problem of multiple realizability. A prosthetic model might be affordance valid and phenomenally valid yet mechanistically invalid. Prosthetic runners' legs do not work like typical biological legs. Heart and lung machines do not work like hearts and lungs. If so, then building a functional prosthesis that simulates a mechanistic model is insufficient to demonstrate that the model is mechanistically valid.

One might proceed from this point to compare the internal states and transitions of the simulation in the prosthesis with the internal states and transitions of the target, but this comparison is no more and no less difficult for prosthetic modeling than it is for simulation.

Next consider the model's phenomenal validity. The crucial difference between prosthetic and nonprosthetic simulations lies in the interfaces. To what extent should success in building a prosthesis increase confidence that one has correctly identified the target's working inputs and outputs?

Judging from current prosthetic devices, the epistemic problem of multiple realizability discussed above arises again at the level of inputs and outputs. On the input side, consider the use of brain-machine-interface (BMI) devices to drive a cursor around a computer screen, to play brain-PONG (a version of the 1972 video game made more exciting by the fact that it is played with electroencephalogram [EEG] waves), and to move a robotic arm (see Schwartz 2004; Lebedev and Nicolelis 2006; Schwartz et al. 2006; Velliste et al. 2008). Brain signals are the inputs to these devices. Behaviors are the outputs. No currently available BMI device, to my knowledge, makes use of just those brain inputs that move limbs in typical animals. EEG interface systems use gross cortical waves; electrophysiological interface systems are driven by 15–200 neurons. The goal, rather,

https://doi.org/10.1086/656822 Published online by Cambridge University Press

is to sample the electrical activity across the cortex to find signals that can be used or that subjects can learn to use to control a prosthesis.

For example, patients can be trained in a few days to use EEG signals to move a robotic arm or a paddle in brain-PONG. EEG signals, however, are widely believed to be epiphenomenal with respect to the mechanisms that produce motor output. They reflect synchronization and desynchronization of cortical activity that correlates with behavior, and one can be trained to voluntarily produce changes in EEG waves to obtain some outcome, but the EEG wave itself is not the functionally significant output of brain motor systems (or so we think). Furthermore, the EEG waves used as input to the prosthesis need not be recorded from the motor cortex: brain waves recorded from auditory cortex, for example, can also be commandeered for this purpose (Felton et al. 2007).

Miguel Nicolelis emphasizes that "precise knowledge of computations performed by brain circuits is not crucial for the construction of clinically relevant BMIs. Mostly BMI platforms take advantage of the well-known correlation between discharges of cortical neurons and motor parameters of interest, and perform a reverse operation: they predict motor parameters from patterns of neuronal firing. Generally, predictions of motor parameters do not signify a causal relationship between the neuronal activity and the generation of movements" (Lebedev and Nicolelis 2006, 540). Indeed, in electrophysiologically driven BMIs, the input signal is typically recorded from 15–200 neurons, which is certainly only an imperfect sample of the neurons that contribute to the behavior in the typical organism, if they contribute at all. The brain states that drive these prostheses need not be states that contribute to movement in subjects with full use of their limbs, and so building a prosthesis need not indicate that one has chosen the correct inputs.

On the output side, the prosthesis must produce outputs that can be used by downstream components in the system. Sensory prostheses must produce outputs that can be interpreted by the nervous system as meaningful signals. However, the lesson from the current generation of sensory prostheses is that the outputs need not be identical to, and in some cases do not even approximate, the signal used by typical biological sensory systems. Recent work on vision substitution, in which camera input is translated into mechanical output on the tongue or back, shows that the output of the prosthesis need not match the typical input to the brain. The substantial capacity of the brain to reorganize in the face of changing input signals further expands the space of affordance-valid outputs from prosthetic devices. The brain is extremely adaptive under the right circumstances. The set of affordance-valid prostheses is arguably much larger than the set of actual mechanisms exhibited in Homo sapiens at the present.

https://doi.org/10.1086/656822 Published online by Cambridge University Press

The lesson to be drawn from these examples is that to interface with the brain is not sufficient to demonstrate that one's model is phenomenally adequate. The space of functional inputs and outputs is larger than the space of functional inputs and outputs that development and evolution have thus far had occasion to exploit. This fact is liberating for engineers: they need not be shackled to the use of inputs and outputs that the standard, untrained, human brain happens to use. They are limited rather by what the brain can reasonably become in the allotted time and in available environments given the ethical, scientific, and technological limits of the age.

The situation is even more complicated for central systems involved in cognitive operations than for peripheral sensory and motor functions. In examples from the periphery, one of the interfaces is sufficiently well characterized that researchers can use it as firm ground from which to explore the other interfaces. In building a prosthetic brain region, both input and output are equally mysterious.

Take the long-range project of Berger and colleagues to build a prosthetic hippocampus. (For a fascinating review of progress to date, see Berger et al. [2005].) The dominant excitatory hippocampal trisynaptic loop projects from dentate gyrus (DG) to the CA3 region and then to the CA1 region. Their first major phase of the project is to create a replacement for the CA3 region. The first step of this phase is to do so in a 400-micrometer hippocampal slice. The prosthesis links DG output to CA1 input in a way that preserves the input-output pattern instantiated by the connections from DG to CA3 to CA1 in the intact rat hippocampus.

Brain regions transform spatiotemporal patterns of activity in a population of presynaptic neurons into spatiotemporal patterns of activity in a population of postsynaptic neurons. Populations of neurons transform the sequence of interspike intervals input to a neuron's dendrites into the sequence of interspike intervals output from the neuron's axon. To build the prosthesis, Berger et al. begin by constructing a mathematical model of the input-output function performed by the CA3 region. They stimulate the DG region in hippocampal slices at physiologically relevant intensities with randomly varying interspike intervals while determining experimentally (i) the interspike intervals and population spike amplitudes in CA3 and (ii) the timing and amplitude of excitatory postsynaptic potentials in CA1. The model constructed on the basis of these experiments is then implemented in a chip that interfaces via a multiunit electrode with both the dentate gyrus and CA1.

Even at this grain of description, it is clear that Berger et al. have chosen inputs and outputs that are different from those used by the typical hippocampus. First, limits of microchip and electrode technology prohibit the authors from modeling and interfacing with each individual neuron

https://doi.org/10.1086/656822 Published online by Cambridge University Press

in the region of interest. Their prosthesis works on population measures of neuronal activity. If this simplifying strategy works (see Berger et al. 2005), the prosthetic model will work on aggregates of the inputs and outputs used by the target mechanism. Second, the authors' current prosthesis works in a razor-thin slice of hippocampus, removed from its context in the rest of the hippocampus and in the brain. The assumption that patterns in individual slices can be added up to produce the corporate behavior of thousands of adjacent slices remains to be tested. Last, in focusing on spike trains, Berger et al. intentionally remove other contextual factors (such as sleep-wake patterns, the endocrine environment of the hippocampus, and any other factors that might vary with aspects of the learning environment, including the content, the training, and other incidental factors).

Berger et al.'s model admirably captures the amplitude of the response in CA1 to trains of impulses delivered to the dentate gyrus. This is remarkable. Yet the authors recognize it is only a first step to demonstrating that such a prosthesis could work. In particular, it does nothing to demonstrate that the DG to CA1 input-output pattern beautifully replicated by the prosthesis is in fact the input-output pattern required for the hippocampus to play its role in memory encoding and storage.

That said, scores of computational models have been proposed for the hippocampus. For none of them is it currently possible to demonstrate that an animal would learn if the hippocampus worked as the model suggests. A successful prosthesis would answer this challenge definitively. To answer this challenge, however, is not to demonstrate that one understands how the target mechanism works. One can build a prosthesis by commandeering signals available in the brain for new purposes rather than by building a device that can take target inputs as prosthetic inputs and give target outputs as prosthetic outputs.

This point is driven home when one considers the well-documented plasticity of the brain (see Nicolelis 2003, 421). Just as brain systems recover and rewire in response to brain damage and disease, they also reorganize to accommodate new devices (Lebedev et al. 2005). This is one of the fundamental lessons driven home by BMI research. For example, subjects can learn reliably to produce motor output linked to the behavior of as few as 15 cortical cells. Indeed, BMI research is now routinely directed at understanding how the brain changes to accommodate these devices.

These considerations show that it is helpful to distinguish phenomenal validity from affordance validity. Phenomenal validity is a matter of whether the model instantiates the input-output relation by which the target interfaces with its environment. Affordance validity, in contrast, is a matter of whether the model instantiates an input-output relationship

https://doi.org/10.1086/656822 Published online by Cambridge University Press

that the brain could use to replace the target's function. A prosthetic model demonstrates that the brain can be made to work *with* a simulation, even if it falls short of demonstrating that the brain works *like* the simulation.

Affordance validity is a matter of degree. A prosthetic with no affordance validity cannot be used to replace or restore function. A prosthetic with maximal affordance validity performs optimally. Qualitatively, a "commandeering index" to reflect affordance validity would depend on (1) the capacity of the target to accept inputs, (2) the capacity of the target to adapt to new inputs, (3) features of the subject's social and physical environment, (4) the available interface technology, and (5) the mechanistic and phenomenal details concerning the implemented model. It remains to be seen whether phenomenally and mechanistically valid interfaces are, in fact, optimal (achieving the highest value of affordance validity) or whether (as seems likely) it is possible to alter nature's handiwork and find new ways to do old things with brains.

4. Explaining and Applying Brain Knowledge. The above considerations emphasize three significant differences between explanatory knowledge of how the brain works and maker's knowledge of how to prevent disease, repair damage, and recover function. The phrase "maker's knowledge" is intended to capture the ideal of scientific knowledge expressed by Bacon in the *Novum Organum* (1620; see Pérez-Ramos 1988).

First, the quest for explanatory knowledge focuses attention on a narrower range of possible mechanisms than the quest for maker's knowledge. To explain a phenomenon, one seeks to understand how a mechanism actually works in a target system; to apply knowledge, one needs to know how the brain can be made to work within the target system. Just as sprinters with prosthetic legs run without mimicking the structure of genetically human legs, one might build a functional prosthetic brain region or interface without mimicking the structure of the target. Explanatory models must be ecologically and/or physiologically valid to an extent that prostheses need not be. Explanatory modelers are limited by their ambitions to describe just the components, activities, and modes of organization actually used in the target (cf. Lebedev and Nicolelis 2006, 540). Prostheses can to some extent presuppose unnatural environments and unusual physiology.

Second, engineers are limited by available technology. The available instruments for building interfaces are limited by spatial and temporal resolution and bandwidth. The choice of technology for use in the prosthesis is influenced in large part by considerations of biocompatibility, cost, durability, invasiveness, and safety. Such ethical, practical, and tech-

https://doi.org/10.1086/656822 Published online by Cambridge University Press

nological limitations are less germane to the explanatory project (see Schwartz et al. 2006, 216).

Finally, the engineer need not respect the target mechanism's shortcomings. To understand the mechanisms underlying normal function, one needs to explain not only the mechanism's typical behavior but also how it behaves in nonstandard conditions, when it is given abnormal inputs, and when it breaks (Craver 2006, 2007). The engineer can ignore such frailties. Avoiding the frailties of memory might well be treated as a virtue of a prosthetic hippocampus rather than as a shortcoming. A biologically realistic model of the hippocampus, in contrast, must respect the illusions of memory, reproduce the learning and forgetting curves, and explain why damage to the hippocampus produces its paradigmatic deficits.

Once these twin aims of explanation and application are before us, it is plain that excessive devotion to ecological and physiological relevance of one's investigations into the mechanisms of the brain limits unnecessarily the search for maker's knowledge of how to make the brain work, to change its physiological organization, and to alter its physical and social environment so as to unlock new powers of the brain. In the search for applications of our knowledge of brain function, prosthetic modelers (and other medical researchers) are free to think outside the box, to imagine new physiologies and new environments that might allow the brain to do new things.

5. Conclusion. Prosthetic models are increasingly prominent in contemporary neuroscience. They are used both for practical applications and as experimental tools. Here, I have asked whether the ability to build a successful prosthesis counts as evidence that one knows how the system works. I have introduced the notion of affordance validity as a measure of whether the model in the prosthesis could be made to work for us in the context of a biological system. I argue that affordance valid models need not be mechanistically or phenomenally valid. This is a blessing for engineers and a mild epistemic curse for basic researchers. However, the very question leads us to reflect on the complex relationship between the twin goals of neuroscience: explanation and control.

REFERENCES

Bacon, F. 1620. *Novum Organum*. London.

Benzania, F. 1997–2003. Nerve (version 2), http://nerve.bsd.uchicago.edu/nerve1.html.

Berger, T. W., et al. 2005. "Restoring Lost Cognitive Function." *IEEE Engineering in Medicine and Biology Magazine* 24 (September/October): 30–44.

Bogen, J. 2005. "Regularities and Causality: Generalizations and Causal Explanations." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:397–420.

Craver, C. F. 2001. "Role Functions, Mechanisms and Hierarchy." *Philosophy of Science* 68:31–55.

https://doi.org/10.1086/656822 Published online by Cambridge University Press

———. 2006. "When Mechanistic Models Explain." *Synthese* 153:355–76.
———. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon.
———. 2008. "Physical Law and Mechanistic Explanation in the Hodgkin and Huxley Model of the Action Potential." *Philosophy of Science* 75:1022–33.
Cummins, R. 1975. "Functional Analysis." *Journal of Philosophy* 72:741–64.
Dayan, P., and L. F. Abbott. 2001. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge, MA: MIT Press.
Destexhe, A., and T. Bal, eds. 2009. *Dynamic-Clamp: From Principles to Applications*. Series in Computational Neuroscience. New York: Springer.
Felton, E. A., A. Wilson, J. C. Williams, and G. P. Charles. 2007. "Electrocorticographically Controlled Brain-Computer Interfaces Using Motor and Sensory Imagery in Patients with Temporary Subdural Electrode Implants: Report of Four Cases." *Journal of Neurosurgery* 106:495–500.
Guala, F. 2002. "Models, Simulations, and Experiments." In *Model-Based Reasoning: Science, Technology, Values*, ed. L. Magnani and N. J. Nersessian, 59–74. New York: Kluwer.
Hartman, S. 1996. "The World as a Process: Simulations in the Natural and Social Sciences." In *Modeling and Simulation in the Social Sciences from the Philosophy of Science Point of View*, ed. R. Hegselmann, U. Müller, and K. Troitzsch, 77–100. Theory and Decision Library. Dordrecht: Kluwer.
Hille, B. 1992. *Ion Channels of Excitable Membranes*. 2nd ed. Sunderland, MA: Sinauer.
Hodgkin, A. L., and A. F. Huxley. 1952. "A Quantitative Description of Membrane Current and Its Application to Conduction and Excitation in Nerve." *Journal of Physiology* 117:500–544.
Korb, K. B., and S. Mascaro. 2009. "The Philosophy of Computer Simulation." In *Logic, Methodology, and Philosophy of Science: Proceedings of the Thirteenth International Congress*, ed. C. Glymour, W. Wang, and D. Westerstahl. London: College Publications.
Lebedev, M. A., J. M. Carmena, J. E. O'Doherty, M. Zacksenhouse, C. S. Henriquez, J. C. Principe, and M. A. L. Nicolelis. 2005. "Cortical Ensemble Adaptation to Represent Actuators Controlled by a Brain Machine Interface." *Journal of Neuroscience* 25:4681–93.
Lebedev, M. A., and M. A. L. Nicolelis. 2006. "Brain Machine Interfaces: Past, Present and Future." *Trends in Neurosciences* 29:536–46.
Marr, D. 1981. "Artificial Intelligence: A Personal View." In *Mind Design*, ed. J. Haugeland, 129–42. Cambridge, MA: MIT Press.
Mauk, M. 2000. "The Potential Effectiveness of Simulations versus Phenomenological Models." *Nature Neuroscience* 3:649–51.
Nicolelis, M. A. L. 2003. "Brain-Machine Interfaces to Restore Motor Function and Probe Neural Circuits." *Nature Reviews Neuroscience* 4:417–22.
Parker, W. 2009. "Does Matter Really Matter? Computer Simulations, Experiments, and Materiality." *Synthese* 169 (3): 483–96.
Pérez-Ramos, A. 1988. *Francis Bacon's Idea of Science and the Maker's Knowledge Tradition*. Oxford: Oxford University Press.
Prinz, A. A., L. Abbott, and E. Marder. 2004. "The Dynamic Clamp Comes of Age." *Trends in Neurosciences* 27:218–24.
Schwartz, A. B. 2004. "Cortical Neural Prostheses." *Annual Review of Neuroscience* 27:487–507.
Schwartz, A. B., X. T. Cui, D. J. Weber, and D. W. Moran. 2006. "Brain-Controlled Interfaces: Movement Restoration with Neural Prosthetics." *Neuron* 52:205–20.
Touretzky, D. S., M. V. Albert, N. D. Daw, A. Ladsariya, and M. Bonakdarpour. 2008. HHsim: Graphical Hodgkin-Huxley Simulator, http://www.cs.cmu.edu/~dst/HHsim/.
Velliste, M., S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz. 2008. "Cortical Control of a Prosthetic Arm for Self-Feeding." *Nature* 453 (June 19): 1098–1101.
Weisberg, M. 2007. "Who Is a Modeler?" *British Journal for Philosophy of Science* 58:207–33.

https://doi.org/10.1086/656822 Published online by Cambridge University Press