Adaptive Behavior

Animals *and* Animats: Why Not Both Iguanas?

Randall D. Beer[1,2,3], Paul L. Williams[1]
[1] *Cognitive Science Program, Indiana University*
[2] *Department of Computer Science, Indiana University*
[3] *Department of Informatics, Indiana University*

1 Introduction

In her target article, Webb contrasts two kinds of models, which she calls animal and animat models, and argues that the latter are unfairly held to less strict standards of scientific relevance than the former, particularly in regards to being subjected to empirical refutation. In order to illustrate her position, she draws upon both her own work on cricket phonotaxis (Reeve & Webb, 2003; Webb, 1995) and our work on the evolution and analysis of model brain–body–environment systems (for a review see Beer, 2008), focusing specifically on our studies of categorical perception (Beer, 2003a). We applaud Webb for engaging the general issue of model interpretation in adaptive behavior and artificial life, since it is far too common for work in these communities to be at best unclear and at worst intentionally ambiguous about their intended scientific relevance. We also completely agree that any scientific research must ultimately be judged by the degree to which it illuminates the actual phenomenon of interest and that, to the extent that animats are scientifically relevant, they are indeed models.

However, we could not disagree more strongly with Webb's overly restrictive conception of the kind of models they must be. Any model can be characterized by its answers to several key questions. What is the model's target? How does the model relate to its target? What purpose is the model intended to serve? How should the model's success be evaluated? We will call these questions the *fundamental modeling questions*. As we understand it, Webb's central argument turns on her insistence that these questions be answered in a particular way, which is grounded in a specific modeling methodology that we will call *data-driven modeling*. What she fails to recognize, however, is that there are many other kinds of modeling methodologies that offer different but equally valid answers to these questions. In particular, our own work is grounded in the tradition of *theory-driven modeling*, which has its roots in physics. In this commentary, we attempt to briefly characterize both data-driven and theory-driven modeling and to contrast the sorts of answers they give to the fundamental modeling questions.

2 Data-Driven Modeling

Webb's argument assumes a recipe for modeling that goes something like this. Through years of painstaking experimental work, biologists collect extensive data about certain aspects of, say, a neuroethological system of interest. The first job of the modeler is to synthesize some subset of this existing data into a model of that system. This model may take different forms, including mathematical, computational, or robotic, depending on the nature of the system being modeled, the particular scientific questions of interest, and the skills of the modeler. Once constructed, experiments can be carried out on the model. For example, a model of cricket phonotaxis can be placed in different acoustic environments, or particular model neurons or connections can be lesioned. If these experiments

Correspondence to: Randall D. Beer, Cognitive Science Program, 1910 E. 10th St. – 840 Eigenmann, Indiana University, Bloomington, IN 47406, USA. *E-mail*: rdbeer@indiana.edu. *Tel.*: (812) 856-0873, *Fax*: (812) 855-1086.

(2009), Vol 17(4): 296–302.
DOI: 10.1177/1059712309340844

have been performed previously on the animal, then they become independent tests of the model's validity. In contrast, if the model is manipulated in ways that have not yet been attempted in the animal, then the results of these manipulations become specific quantitative predictions that can be tested by subsequent experiments. The results of such experiments may either serve to confirm or refute the model, triggering another cycle of model-building and testing.

The virtues of the model–prediction–test–refinement loop of data-driven modeling are many and well known. The very process of constructing such a model can be a useful exercise because it can reveal important holes in what the experimentalists may have thought was a fairly complete data set, leading to additional experiments to fill these holes. Once an initial model is constructed, it can be used to test the functional adequacy of prior understanding of the system. Quite often, this understanding will be shown to be inadequate because of complex nonlinear interactions between various components of the system. Indeed, the ability to explore such interactions is a significant advantage of modeling. Once a model has been sufficiently verified, it can also serve as a kind of stand-in for the original biological system, supporting observations or manipulations that are difficult or impossible in vivo, at least with current technology. Thus, over time, this process generates a powerful feedback loop between experiment and modeling, with the model becoming a kind of interactive repository of the best current data and understanding of a biological system.

For data-driven modeling, the ways in which the fundamental modeling questions should be answered are clear. The target for a model should be some specific biological system (not literally a single biological individual, but rather a set of closely related individuals, such as a species). Moreover, the model should relate to its target by serving as a more or less accurate representation of it. Note that the required accuracy depends on the biological question of interest. For example, if one is interested in an orientation behavior such as phonotaxis, then one might hope that the details of the ionic channels in the underlying nerve cells or the particular means of locomotion utilized are not essential (although the low-level details of biological systems sometimes have a tendency to penetrate to higher levels). Data-driven modeling can be carried out for a variety of purposes, some of which were described above. However, regardless of the purpose that the model is intended to serve, the standards for evaluation are always the same: validation through empirical testing. If experiments consistently fail to verify a model prediction, then the model is refuted and must be revised accordingly. Furthermore, given the reductionist tendencies of biology, experimentalists typically ask increasingly more detailed questions about a system over time, requiring a data-driven model to incorporate more and more realistic detail in order to remain relevant.

The data-driven methodology is a completely traditional perspective on modeling, typical of experiment-dominated sciences such as biology. This is a game whose rules we understand very well, having played it ourselves for many years (e.g., Beer, Kacmarcik, Ritzmann, & Chiel, 1991; Sutton et al., 2004; Synder, Sutton, Neustadter, Beer, & Chiel, 2006). However, by attempting to impose the criteria of data-driven modeling on all scientific models, Webb goes too far. Data-driven modeling is not, as they say, the only game in town. It is not even the best game in town for certain kinds of scientific activities. Different games have different rules, and different means of keeping score. In the next section, we describe another modeling methodology, theory-driven modeling, which provides distinctive but equally valid answers to the fundamental modeling questions.

3 Theory-Driven Modeling in Physics

"Consider a spherical cow…" serves as the punch line of countless jokes about physicists being asked to improve milk production by a dairy farmer. The target of such jokes is the (theoretical) physicist's tendency to reduce a problem to its simplest possible form in order to begin to reason about its solution. What makes such jokes funny is that treating a cow as a sphere seems like such a ridiculous distortion of the experimental facts that no scientifically useful conclusions could possibly follow from it. And yet, as Krauss (1993) has convincingly argued, remarkable insight can often be gained by just such extreme simplifications. For example, it helps explain why simply making cows bigger will not work because, since surface area scales quadratically while volume scales cubically, the internal pressure of increasingly larger cows would eventually cause their skin to burst. Even more to the point, small extensions to this analysis can explain why dinosaurs had to have such small heads relative to their bodies while whales do not, or why the balance of thermal and gravitational pressure that stabilizes stars takes the form that it does, since this reasoning depends only on very general geometrical features of physical objects in three-dimensional space.

The paradigmatic example of this kind of reasoning is provided by Galileo, who was able to uncover the deeper laws of motion by stripping away the bewildering variety of its surface manifestations. Another example is the Ising model, which has played a fundamental role in understanding phase transitions in solid state systems such as ferromagnets (which spontaneously magnetize in a certain temperature range), not because it models the specifics of such physical systems particularly accurately, but because it exhibits some of the same general phenomenology yet is more amenable to analysis (Brush, 1967). The Ising model has since found wide application outside of solid state physics, including as a model of neural networks, flocking birds, or beating heart cells, and has even found applications in economics and sociology. Likewise, ongoing attempts to reconcile quantum mechanics (the theory of the very small) and general relativity (the theory of the very massive or energetic) in situations where both apply (such as inside black holes and near the time of the big bang) have relied heavily on theory-driven models:

> Attempts to reconcile quantum theory and general relativity date back to the 1930s, but despite decades of hard work, no one has succeeded in formulating a complete, self-consistent quantum theory of gravity.… The obstacles to quantizing gravity are in part technical.… But the problem of finding a consistent quantum theory of gravity goes deeper. General relativity is a geometric theory of spacetime, and quantizing gravity means quantizing spacetime itself. In a very basic sense, we do not know what this means.
>
> Faced with such problems, it is natural to look for simpler models that share the important conceptual features of general relativity … General relativity in 2+1 dimensions – two dimensions of space and one of time – is one such model.… With a few exceptions, (2+1)-dimensional solutions are physically quite different from those in 3+1 dimensions, and the (2+1)-dimensional model is not very helpful for understanding the dynamics of realistic quantum gravity. But for the analysis of conceptual problems – the nature of time, the construction of states and observables, the role of topology and topology change, the relationship among different approaches to quantization – the model has proven highly instructive. (Carlip, 1998).

This kind of reasoning is ubiquitous in theoretical physics, and it leads to a theory-driven modeling methodology whose answers to the fundamental modeling questions are quite different from those of the data-driven methodology (although data-driven modeling is obviously also employed heavily in physics). These models (often called *simplified*, *toy*, or *idealized* models in the literature) are not grounded in the specifics of any particular existing natural system. Rather, their target is the set of such systems that exhibit some phenomena of interest, and they relate to their targets by attempting to capture the essential features that are common to all members of this set. The purpose of a theory-driven model is to explore some fundamental conceptual issues raised by this class of systems. Finally, theory-driven models are judged by their ability to illuminate the technical or conceptual problems that motivate them. Although such models are sometimes used to make predictions (typically of a qualitative rather than quantitative form), this is not their central purpose and to criticize such models for their lack of empirical grounding is to miss their very point. They are tools for thinking. They are not only intuition pumps, but theory pumps, and even experiment pumps. They are the bread and butter of the theoretical development of any science.

4 Theory-Driven Modeling in Biology

Lest we give the impression that theory-driven modeling is unique to physics, let us briefly describe four examples from biology. First, consider the Fitzhugh–Nagumo model of neural excitability (Fitzhugh, 1961). As Webb mentions, the Hodgkin–Huxley model is a famous data-driven model of the generation of action potentials in the squid giant axon whose extensions have found wide application throughout neuroscience. However, what she fails to mention is that, although the Hodgkin–Huxley model can accurately fit the electrical behavior of many different nerve cell types in many different species, its dynamical properties have been very difficult to understand. For example, the subtle nature of the threshold in the Hodgkin–

Huxley model was only illuminated by studying idealized models, such as the Fitzhugh–Nagumo model, which have no quantitative grounding in electrophysiology (Izhikevich, 2007).

Second, consider Boolean network models of genetic regulatory networks (Kauffman, 1969), in which genes are idealized as binary switches and regulatory interactions are idealized as Boolean functions. These models have not only been applied in a data-driven manner to the analysis of specific regulatory networks, such as Drosophila's segment polarity network (Albert & Othmer, 2003), but have also been used to study the properties of regulatory networks in general, such as robustness, modularity, and evolvability (Ciliberti, Martin, & Wagner, 2007; Fernández & Solé, 2006). Third, consider Hinton and Nowlan's (1987) simple model of the Baldwin effect, in which changes occurring during the lifetimes of individuals can become genetically assimilated into their species, speeding up evolution in the process. Although it is not grounded in any particular animal system, this model has played a crucial role in demonstrating the reality of the Baldwin effect and stimulating its further study (Maynard Smith, 1987). Finally, in some of our own recent work, a simplified model of the development of neural circuits has been used to explore the origin, nature, and consequences of developmental bias in evolution (Psujek & Beer, 2008).

It was in the spirit of this long tradition of theory-driven modeling that our own "toy" neuroethological models were developed. To date, these models have targeted issues in sensorimotor behavior, learning, and cognition (for a review see Beer, 2008). From the outset, the overall goal of our minimally cognitive behavior work was to investigate the cognitive implications of situatedness, embodiment, and dynamics using the kinds of tasks that cognitive scientists themselves study (Beer, 1996, 2003a). Our specific interest in categorical perception was motivated by the claim that categorization is a foundational component of cognitive processes (Harnad, 1987).

Since categorical behavior is common throughout the animal kingdom, what might be called the "proximal" target of our model was the set of general features that are common across manifestations of categorical behavior. Given the cognitive focus of this work, we were especially interested in the more sophisticated forms of categorical perception, such as object categorization and, more recently, relational categorization (Williams, Beer, & Gasser, 2008). Accordingly, the form of our model was derived from a principled simplification of key features of object categorization. Since object categorization usually involves vision, we chose to employ an array of distal sensors as a simple model of vision. Given our interest in embodiment, we chose to deploy these sensors on a body that was capable of the simplest possible kind of motion, movement along one dimension. For the same reason, we also chose to have the agent express its decision through approach or avoidance behavior rather than simply through an abstract category signal. Finally, we chose to utilize continuous-time recurrent neural networks because, as we have argued repeatedly elsewhere, they are the simplest neural network model that captures the nonlinear dynamical properties of nervous systems. Note that other features of our model, such as the fact that the agent has seven ray sensors each of whose maximum length is 220, are completely arbitrary. The arbitrariness of still other features, such as the particular connection weights evolved in one run of the evolutionary algorithm, is mediated by the fact that the evolutionary algorithm can be run many times, resulting in a set of effective categorizers with different neural parameters. Indeed, our evolutionary approach to theory-driven modeling is not really designed to produce a single model, but rather a family of models whose common properties can then be analyzed.

However, what might be called the "distal" target of our model agent was not really categorical perception at all, but rather the dynamical analysis of the simplest possible model agents exhibiting situated and embodied cognitive behavior. As stated explicitly in Beer (2003b), the purpose of the model was threefold: (a) "to encourage a direct confrontation of situated, embodied and dynamical ideas with cognitive phenomena, in the hopes of furthering debate on the roles of these ideas in cognitive science"; (b) "to describe a specific research methodology that could concretely ground such a debate, and to illustrate in some technical detail how a situated, embodied, minimally-cognitive agent could be developed and dynamically analyzed"; (c) "to use this model agent as a springboard to begin to explore some of the larger implications of these ideas for explanation in cognitive science". The reason for performing a psychophysical analysis of the model was simply to demonstrate that it could be approached in much the same way as a natural instance of categorical perception, not to make dubious claims about the model's psychological or neurophysiological validity. Ultimately, like any theory-driven model, this model must be judged by the extent to which it illuminates the conceptual issues it set out to explore. But regardless of how successful it turns out to have been in this regard, or how one feels about the particular modeling choices that were made, we do not see how this model differs in kind from, for example, the "made-up" (2+1)-dimensional model universes that are studied in quantum gravity in order to grapple with the conceptual issues arising in the quantization of space-time.

Despite the fact that the primary purpose of such toy models is the exploration of general conceptual issues, can they ever contribute specific insights to empirically grounded studies of particular animals? Indeed, they can, and we end this section with two examples from our own work. The first example, which we mention only briefly, involved a comparison of evolved model agents for selective attention (Slocum, Downey, & Beer, 2000) with human subjects on the same task, demonstrating a common reliance on reactive inhibition (Ward & Ward, 2008). The second example concerns our work on the evolution and analysis of model pattern generators for walking (Beer, Chiel, & Gallagher, 1999; Beer & Gallagher, 1992; Chiel, Beer, & Gallagher, 1999). The goal of this work was to explore the conditions under which different pattern-generation architectures evolve, the interplay between neural and mechanical properties in sensorimotor behavior, and the extent to which central pattern generators (CPGs) can be modularly decomposed.

In analyzing a population of one hundred evolved highly fit CPGs, we discovered that, although we could identify common principles at an appropriately abstract level of analysis using dynamical systems theory, the evolved neural parameters themselves varied so much that simply averaging them produced a network that failed to even oscillate. We then pointed out that this finding could have important experimental implications, since neurobiologists generally average the results of experiments on multiple animals in order to obtain parameter values for biophysically realistic neural models. Our results led Golowasch, Goldman, Abbott, and Marder (2002) to examine the same issue in models of the stomatogastric ganglion, a CPG that controls a completely different behavior (the chewing motions of cartilaginous teeth in the stomach) of a completely different animal (the lobster). They verified that exactly the same problem occurs in more biophysically realistic and empirically grounded models. Furthermore, our discovery of multiple instantiability and sensitivity and robustness to parameter perturbations in our evolved model agents anticipated the discovery of similar phenomena in the stomatogastric ganglion (Goldman, Golowasch, Marder, & Abbott, 2001; Prinz, Bucher, & Marder, 2004).

5 Conclusion

Empirically grounded models play an absolutely vital role in science. But science is fundamentally about understanding the natural world, not just accounting for the results of specific experiments. Our distinction between data-driven and theory-driven models aligns closely with the practical/theoretical distinction made in a recent textbook on biological modeling (Ellner & Guckenheimer, 2006). While data-driven models attempt to obtain general insights through the study of specific systems, theory-driven models attempt to study general issues directly by formulating highly simplified models that focus on the commonalities of many specific systems. Both kinds of models are essential to science, and concerns have repeatedly been raised about the dangers of emphasizing data-driven modeling over theory in biology (De Schutter, 2008; Lazebnik, 2004; McCollum, 2000).

In fact, models can contribute to understanding in many other ways beyond those considered here. For example, Marder and Abbott (1995) recognize three different kinds of models: speculative, confirmatory, and interpretative. In a recent essay, Epstein (2008) identifies seventeen different uses of models, only one of which is prediction and fewer than half of which have anything to do with empirical grounding. The relationship between these various kinds of models is not competitive, but complementary. There is no question that modelers should be as clear as possible about the nature of the intended relevance of their models. However, we find it more than a little ironic that Webb is arguing to restrict the role of models in biology at the same time that biologists themselves are increasingly turning to the theoretical approaches of physics and mathematics to help them better understand the daunting complexity of living systems (Cohen, 2004).

References

Albert, R., & Othmer, H. G. (2003). The topology of the regulatory interactions predicts the expression pattern of the segment polarity genes in *Drosophila melanogaster*. *Journal Theoretical Biology*, *223*, 1–18.

Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes, M. Mataric, J. Meyer, J. Pollack, & S. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 421–429). Cambridge, MA: MIT Press.

Beer, R. D. (2003a). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, *11*, 209–243.

Beer, R. D. (2003b). Arches and stones in cognitive architecture. *Adaptive Behavior*, *11*, 299–305.

Beer, R. D. (2008). The dynamics of brain-body-environment systems: A status report. In P. Calvo & A. Gomila (Eds.), *Handbook of cognitive science: An embodied approach* (pp. 99–120). Elsevier.

Beer, R. D., Chiel, H. J., & Gallagher, J. C. (1999). Evolution and analysis of model CPGs for walking II. General principles and individual variability. *Journal of Computational Neuroscience* , *7*(2), 119–147.

Beer, R. D., Kacmarcik, G. J., Ritzmann, R. E., & Chiel, H. J. (1991). A model of distributed sensorimotor control in the cockroach escape turn. In R. P. Lippman, J. Moody, & D. S. Touretzky (Eds.), *Advances in neural information processing systems 3* (pp. 507–513). Morgan Kaufmann.

Brush, S. G. (1967). History of the Lenz-Ising model. *Reviews of Modern Physics*, *39*, 883–893.

Carlip, S. (1998). *Quantum gravity in 2+1 dimensions*. Cambridge University Press.

Chiel, H. J., Beer, R. D., & Gallagher, J. C. (1999). Evolution and analysis of model CPGs for walking I. Dynamical modules. *Journal of Computational Neuroscience*, *7*(2), 99–118.

Ciliberti, S., Martin, O. C., & Wagner, A. (2007). Robustness can evolve gradually in complex regulatory gene networks with varying topology. *PLoS Computational Biology*, *3*(2), e15. DOI:10.1371/journal.pcbi.0030015

Cohen, J. E. (2004). Mathematics is biology's next microscope, only better; biology is mathematics' next physics, only better. *PLoS Biology*, *2*(12), e439. DOI:10.1371/ journal.pbio.0020439

De Schutter, E. (2008). Reviewing multi-disciplinary papers: A challenge in neuroscience? *Neuroinformatics*, *6*, 253–255.

Ellner, S. P., & Guckenheimer, J. (2006). *Dynamical models in biology*. Princeton, NJ: Princeton University Press.

Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation 11*(4), 12. Retrieved from http://jasss.soc.surrey.ac.uk/11/4/12.html

Fernández, P., & Solé, R. V. (2006). The role of computation in complex regulatory networks. In E. V. Koonin, Y. I. Wolf, & G. P. Karev (Eds.), *Power laws, scale-free networks and genome biology* (pp. 206–225). Berlin: Springer.

Fitzhugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, *1*, 445–466.

Goldman, M. S., Golowasch, J., Marder, E., & Abbott, L. F. (2001). Global structure, robustness, and modulation of neuronal models. *Journal of Neuroscience*, *21*(14), 5229–5238.

Golowasch, J., Goldman, M. S., Abbott, L. F., & Marder, E. (2002). Failure of averaging in the construction of a conductance-based neuron model. *Journal of Neurophysiology*, *87*, 1129–1131.

Harnad, S. (Ed.) (1987). *Categorical perception: The groundwork of cognition*. Cambridge University Press.

Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, *1*, 495–502.

Izhikevich, E. M. (2007). *Dynamical systems in neuroscience*. Cambridge, MA: MIT Press.

Kauffman, S. A. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, *22*, 437–467.

Krauss, L. M. (1993). *Fear of physics*. Basic Books.

Lazebnik, Y. (2004). Can a biologist fix a radio? – Or, what I learned while studying apoptosis. *Cancer Cell*, *2*, 179–182.

Marder, E., & Abbott, L. F. (1995). Theory in motion. *Current Biology*, *5*, 832–840.

Maynard Smith, J. (1987). When learning guides evolution. *Nature*, *329*, 761–762.

McCollum, G. (2000). Social barriers to a theoretical neuroscience. *Trends in Neurosciences*, *23*, 334–336.

Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, *7*(12), 1345–1352.

Psujek, S., & Beer, R. D. (2008). Developmental bias in evolution: Evolutionary accessibility of phenotypes in a model evo-devo system. *Evolution and Development*, *10*, 375–390.

Reeve, R. E., & Webb, B. H. (2003). New neural circuits for robot phonotaxis. *Philosophical Transactions of the Royal Society of London A*, *361*, 2245–2266.

Slocum, A. C., Downey, D. C., & Beer, R. D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, & S. Wilson (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior* (pp. 430–439). Cambridge, MA: MIT Press.

Sutton, G. P., Mangan, E. V., Neustadter, D. M., Beer, R. D., Crago, P. E., & Chiel, H. J. (2004). Neural control exploits changing mechanical advantage and context dependence to generate different feeding responses in *Aplysia*. *Biological Cybernetics*, *91*, 333–345.

Synder, V. A., Sutton, G. P., Neustadter, D. M., Beer, R. D., & Chiel, H. J. (2006). Mechanical reconfiguration mediates swallowing and rejection in *Aplysia californica*. *Journal of Comparative Physiology A*, *192*, 857–870.

Ward, R., & Ward, R. (2008). Selective attention and the control of action: Comparative psychology of an artificial evolved agent and people. *Journal of Experimental Psychology*, *34*(5), 1165–1182.

Webb, B. H. (1995). Using robots to model animals: A cricket test. *Robotics and Autonomous Systems*, *16*, 117–134.

Williams, P. L., Beer, R. D., & Gasser, M. (2008). An embodied dynamical approach to relational categorization. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 223–228).