First, there is a biological component to the workings of political institutions because they are "inhabited" by human beings. Thus, the functioning of a political institution depends at least in part on how it interacts with the human body and brain. Trivially, an institution that does not respect the human need for sleep is not going to last for long. Less trivially, well-functioning institutions respect the quirks of human cognition, emotion, and moral reasoning. They rely on formal structures (simple majority rule, federalism, written constitutions) and informal modes of interaction (culture, ideology, morale, and leadership). They take into account limitations on the human ability to pay attention (what do people naturally attend to? what is the human attention span?). They pander to the human desire for narrative (people tend to work hard for institutions that have a vision and give them a sense of purpose). Great institutions make productive use of the human propensity for factionalism (people tend to form clusters, and they cooperate within and compete across clusters; people also seem to have a close-to-infinite capacity to form clusters-within-clusters).

Second, political institutions are complex systems with emergent properties. Just like the brain possesses consciousness and yet one cannot find the seat of consciousness in it, a political institution can contribute to economic stability and growth even though its internal world consists of a Kafkaesque bureaucracy obsessed with partisan squabbles and prone to petty corruption. We have theories connecting the internal world of an institution to its aggregate performance. Even so, there is something magical about well-performing institutions, which is to say that our theories are thin simplifications.

Third, political institutions evolve in conflict with their environment. Some institutions emerge bottom up as a result of the decentralized actions of individuals. In a process comparable to annealing, people's interactions become increasingly structured until a complex institution crystallizes out of the mess.[3] Other institutions are designed from scratch by political actors, but they evolve over time as they accommodate the political pressures of disaffected individuals and groups or change their structure to better cope with environmental pressures – or they break down and are replaced by other institutions that are better able to stand the political heat.

So what characterizes a good model of political institutions? For starters, we need to understand why we need a model in the first place. Why can't we just "see and understand" an institution the way it really works? After all, the human brain is capable of seeing and understanding very complex phenomena. Indeed, in many cases (recognizing faces, making sense of gossip) we only realized how complex these phenomena are after our man-made computers choked over analyzing them.

This is where evolutionary psychology comes in. The social nature of the human brain was shaped in a hunter-gatherer environment, roughly 30,000 to 300,000 years ago. People lived together in small groups and spent a lot of time gossiping with and about each other. It is not all that surprising that human beings are good at extracting subtle cues from facial expressions or reading other minds ("if I say this she will think he left unsaid that . . . "). Fortunately or not, there was nothing even remotely resembling the United States Congress floating around in the hunter-gatherer environment. As a result, our brain is poorly equipped to reason about the workings of the United States Congress (or bat sonar for that matter).

The purpose of a model is to take a complex system that our brain cannot comprehend, boil it down to its essential features, and thereby make it transparent. This is where we run into a snag. The underbrush we omit as non-essential is in fact essential to the workings of the system. The thin simplification that is our model would drop down dead if it were forced to confront the environment in which the true complex system survives and thrives. What a good model does is to reconcile as best as it can an irreconcilable tension – it accommodates both the workings of the system (it provides a decent approximation as measured by some yardstick of usefulness) and the cognitive makeup of the human observer (it is illuminating to us).

When it comes to the cognitive makeup of the human observer, there are human universals and then there are cultural and individual-specific differences. Cultural differences include differences in "seeing and believing" across scientific disciplines and subfields – indeed, the process of becoming an expert through graduate training and mentorship can be thought of as a process of getting brainwashed into slicing up the world in a certain way and being blind (and, interestingly, hostile) to other ways of slicing up the world. As a result, different disciplines and subfields come up with different models of complex phenomena.

Indeed, the lack of disagreement carries over to meta-level debate about the purpose and workings of the scientific enterprise. When scientists self-reflect, they come up with competing and partially contradictory models of "what is a model" and "what makes for a good model." And this is good so. If scientific progress occurs, it is because the scientific enterprise is in sync with the cognitive makeup of its human inhabitants – their diversity and their propensity for forming factions and factions-within-factions.

NOTES

1. The ideas developed in this commentary are drawn from the author's published work (e.g., Lohmann 2000).

2. In a given district, the candidate who gets the most votes becomes Member of Parliament regardless of whether he or she is supported by more than 50% of the voters. In the House of Commons, the party with the most seats forms the government regardless of whether it has won a majority of seats.

3. The concept of annealing comes out of the thermodynamics of how liquids freeze and crystallize or metals cool and anneal. At high temperatures, the molecules of a liquid metal move around freely. If the liquid is cooled slowly, the atoms line up to form a pure crystal that is completely regular and represents the state of minimum energy for the system. If the liquid is cooled quickly, it does not reach the minimum energy state but instead gets stuck in a higher energy state. A process that allows the molecules to move around and then gradually quenches them has certain optimality properties.

How building physical models can reduce and guide the abstraction of nature

Malcolm A. MacIver
Division of Engineering and Applied Science & Computation and Neural Systems Program, 104–44, California Institute of Technology, Pasadena, CA 91125. **maciver@robotics.caltech.edu**
http://robotics.caltech.edu/~maciver

Abstract: Animals detect and acquire resources through a sequence of shape changes. This process is tightly coupled to the sensory and mechanical ecology of the animal. Building physical models allow us to prescind from modeling these aspects of the environment, which may not yet be described or suitably abstracted. The significance of this hybrid of physical modeling and experimentation to the acquisition of scientific knowledge is discussed.

The sparsity of resources within a mobile animal's domain compels a certain logic, one that all energy-consuming autonomous agents must follow. Resources must be detected, and behavioral programs engaged that terminate in the acquisition of those resources. In carrying out this imperative of their continued existence, animals exert changes in their geometric configuration in space (in brief, shape) to several ends. These include: (1) for sensing: shape changing to enhance the quality of information from its sensor arrays (e.g., bats manipulating their pinnae position during echolocation; dogs bringing their snouts to the substrate to follow a trail); (2) for locomotion: shape changing to undergo net movement, typically toward the detected resource (e.g., a fish bending its body to swim forward; a biped extending a leg to walk forward);

3) for physical coupling: shape changing to eat or grasp the resource (e.g., depression of the lower jaw of fish to create negative buccal pressure for prey capture; bats flipping the tail membrane up to bring an insect to their mouth).

Animals exhibit an astonishing sophistication in their manipulation of the mechanical properties of their world to achieve these ends. For example, shape changing for sensing in electric fish can be seen in rolling behavior following prey detection (MacIver et al. 2001). This behavior centers the fish's top edge – a region of high sensor density – under the prey; allows the fish to approach the prey by slicing its narrowest cross-section through the water, thereby minimizing added-mass effects; and may provide a simple control strategy for reaching the prey by balancing the stimulus between the two sides of its body and ascending the gradient of sensory signal strength (MacIver et al. 2001; Nelson & MacIver 1999). As described further below, investigations of shape change for locomotion in insects and fish are demonstrating that these animals utilize phenomena within fluids quite beyond those that we utilize in our flying and underwater machines, phenomena that we are still discovering, to say nothing of having a good analytical approach toward.

Shape changing for resource detection and acquisition is clearly fundamental to the sensorimotor intelligence of animals that we so desire to understand. As the examples indicate, these shape changes are tightly coupled to the sensory and mechanical ecology of the animal. Yet, modeling the environment, which animals have demonstrated an unerring capacity to exploit in ways we are hardly aware of, let alone capable of simulating accurately, presents a high obstacle to the integrative computer simulations that are currently our best shot at understanding these coupled sensorimotor processes.

As Webb and others have pointed out (target article, sects. 3.7 & 4.7; Beckers et al. 1996; Flynn & Brooks 1989; Quinn & Espenschied 1993), a great advantage of building physical models is that this allows us to prescind from modeling the undiscovered or unabstracted aspects of the environment on which the target behavior depends. Although Webb's article is very helpful in clarifying the maze of issues surrounding the building of physical models, I believe that this key point is one which merits further elaboration. In what follows, I place the building of physical models in the broader context of the acquisition of scientific knowledge, inquire into the nature of their contribution to this process, and briefly describe some recent examples.

Understanding involves abstraction. These abstractions are expressed in some language for communication and verification. Mathematics provides one such language, but what follows applies to abstractions expressed in any language. Suppose we express our abstractions of some biological phenomenon using the language of mathematics. The next logical step is to calculate predictions from these expressions in order to test their fidelity to the phenomenon (in the case of a spoken language, we would use practices of informal logic to derive verbal predictions). The expressions may need to be approximated to make them computable in finite memory machines in finite time. The calculated predictions are compared to empirically obtained observations, and an interwoven process of theory adjustment, algorithm development, and experimental work ensues. Where can building physical models contribute?

The tragedy of abstraction is that it requires the loss of information. Otherwise, we haven't abstracted. In the process of generating predictions from abstractions, there will be some predictions that will therefore not be computed; namely, those that rely upon the information excluded from our abstractions, or lost in the approximations of those abstractions required by computational expediency. I will use the phrase "abstraction load" (in analogy to "cognitive load") to refer to the work needed to obtain the abstractions and computational methods that will generate the observations we have failed to compute.

Building physical models has the advantage of reducing the potentially insurmountable abstraction load associated with computing all the aspects of the environment on which the target phenomenon depends (where "environment" refers to any aspect external to the phenomenon we are trying to abstract). To simulate the phenomenon adequately, this work would have to be done; but building the object and letting reality supply the physics obviates the need to do some of this work. The crucial point is, we haven't thereby given up the game completely – we are neither pinned into the muck and goo of pure experimentation, nor caged in the assumption-permeated world of pure simulation, but find ourselves at some interesting halfway point.

For example, following similar work by McGeer, Ruina and colleagues developed computational models of a "passive walker" – a walker that has a human-like bipedal gait down inclined planes without actuation or control. The simulations predicted that the walker would not be stable, but it was built in order to test some other issues. To their surprise, the model did walk (Coleman & Ruina 1998). The functioning of the physical model directed the development of a simple quantitative model to explain its stability (Coleman et al. 2001). Similarly, in recent work on fish swimming and insect flying, a number of fluid phenomena have either been discovered or made more observable as a result of the use of robotic devices that approximate the movement of these animals (Ahlborn et al. 1997; Bandyopadhyay et al. 2000; Barrett et al. 1999; Birch & Dickinson 2001; Dickinson et al. 1999).

In allowing the full complexity of the environment to work on what could be called "reduced robotic preparations," this research is cracking open the black box of complex deformable-body and fluid dynamics phenomena to new theoretical advances. The epistemic accessibility afforded by building these robotic devices is analogous to that obtained by traditional instruments such as the microscope and telescope. The building of physical models not only reduces abstraction load, but in illuminating that part of nature we most urgently need to abstract in order to account for a phenomenon, it provides a saliency filter for the immense richness of opportunities for abstraction effort that arise at every turn in the course of experimental work.

When robots fail: The complex processes of learning and development

Ludovic Marin[a] and Olivier Oullier[b]

[a]*Sport, Performance and Health Laboratory, University of Montpellier 1, F-34000 Montpellier, France;* [b]*Center for Complex Systems and Brain Sciences, Florida Atlantic University, Boca Raton, FL 33431.*
l.marin@staps.univ-montp1.fr oullier@walt.ccs.fau.edu
http://oullier.free.fr

Abstract: Although robots can contribute to the understanding of biological behavior, they fail to model the processes by which humans cope with their environment. Both development and learning are characterized by complex relationships that require constant modification. Given present technology, robots can only model behaviors in specific situations and during discrete stages. Robots cannot master the complex relationships that are the hallmark of human behavior.

In her article, Webb offers a convincing argument for instances in which robots can be good models for understanding biological behavior. In her account, mechanical models, like those used in research on other animals, can similarly help researchers gain insight into human behavior. Robots offer experimental advantages in certain situations because robots can be programmed, and even demonstrate very simple learning strategies within a given environment and context. For example, robots can be used as "stand-ins" for humans in experimental situations that are too dangerous for live subjects (e.g., removing the primer of a rocket). Using robots eliminates emotions like fear or anxiety that affect experi-