

obstacle avoidance, and tracking. Faloutsos et al. (2001) show a highly realistic physics-based simulation of complex human body motions such as falling and standing up using a detailed anatomical model of the human skeleton. It would be a major challenge to achieve this degree of anatomical accuracy in a physical robot model.

Even if both implementations are feasible, computer simulations could be more useful tools in modeling biological behavior than robots, since they provide full control over the entire action-perception cycle. Furthermore, simulations are not restricted to real time or real size, so they can represent biological processes that are too slow, too fast, too large or too small for a real-world robot implementation. Neumann and Bühlhoff (2001) use computer simulations to demonstrate that three-dimensional flight with all six degrees of freedom can be visually stabilized using models of spatial orientation strategies found in insects. These strategies exploit the distribution of local light intensities and local image motion in an omnidirectional field of view, and include mechanisms for attitude control, course stabilization, obstacle avoidance, and altitude control. The motor system and flight dynamics of the artificial agent is a simplified model of the fruitfly *Drosophila* and includes effects of drag due to air viscosity. With computer simulations it is possible to represent such effects, which would be extremely difficult to achieve in a robot implementation.

Research, robots, and reality: A statement on current trends in biorobotics

Ernst Niebur,^a Mounya Elhilali,^b Iyad Obeid,^c Justin Werfel,^d Mark Blanchard,^e Mattia Frasca,^f Kaushik Ghose,^g Constanze Hofstoetter,^e Giovanni Indiveri^h and Mark W. Tildenⁱ

^aZanvyl Krieger Mind/Brain Institute and Department of Neuroscience, Johns Hopkins University, Baltimore, MD 21218; ^bInstitute for Systems Research, Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742; ^cDepartment of Biomedical Engineering, Duke University, Durham, NC 27708; ^dDepartment of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139; ^eInstitute of Neuroinformatics, University of Zurich/ETH Zurich, CH-8057 Zurich, Switzerland; ^fDepartment of Electrical, Electronical and Systems Engineering, University of Catania, 95125 Catania, Italy; ^gDepartment of Psychology, University of Maryland at College Park, College Park, MD 20742; ^hFhG-AiS, Fraunhofer Institute for Autonomous Intelligent Systems Technology, 53754 Sankt Augustin, Germany; ⁱInstitute for Physical Sciences, Los Alamos, NM 87544. niebur@jhu.edu <http://cnslab.mb.jhu.edu>
mounya@eng.umd.edu <http://www.ece.umd.edu> io@duke.edu
<http://www.duke.edu/~io> jkwerfel@mit.edu
<http://hebb.mit.edu/people/jkwerfel> jmb@ini.phys.ethz.ch
<http://www.ini.unizh.ch> mfrasca@dees.unict.it
kghose@wam.umd.edu
<http://www.wam.umd.edu/~kghose/research.html>
connie@ini.phys.ethz.ch <http://www.ini.unizh.ch>
giovanni.indiveri@ais.fraunhofer.de <http://www.ais.fraunhofer.de>
mwtilden@qwest.net

Abstract: While robotics has benefited from inspiration gained from biology, the opposite is not the case: there are few if any cases in which robotic models have led to genuine insight into biology. We analyze the reasons why biorobotics has been essentially a one-way street. We argue that the development of better tools is essential for progress in this field.

We will here use the term robot to describe a hardware model of a biological system whose interaction with the physical environment, both in terms of sensors and of actuators, forms an essential part of the model. The question asked in the title of Webb's article is whether such robots can be useful for understanding biology. Our perspective as a group working at the interface between biology and robotics is that robot models have the potential to make considerable contributions, with significant advantages over

other styles of analysis, but that this potential is not being fully exploited at this time.

Robots vs. other models. We compare robot models with computational models (i.e., numerical simulations) and theoretical models (i.e., mathematical abstractions) on the one hand, and biological models, on the other. Computational and theoretical models are devoid of any physical substrate. Although they are comparatively easy to implement, compared to a robotic model of the same biological system, oversimplification due to abstraction gone too far is a significant risk for them. For instance, the modeller must decide what (if any) external noise is to be included and what form to give it; this decision may influence the outcome strongly. A robot model will by its nature be subject to all the actual constraints and conditions of the real world, which cannot be ignored or finessed away. Another disadvantage of computational models is that some properties of the system or its environment may actually be more difficult or costly to simulate in software than in hardware (e.g., nonlinear friction, requirement for real-time response, etc.).

Biological models – those using organisms, cultured cells, brain slices, and so on as their substrate – have other limitations. First, they are vastly more complicated than hardware models, involving complex biological tissue or even whole organisms. Gaining a deep understanding of the system may therefore be difficult. Not only is a robot model simpler than an animal model; since we construct it ourselves, its components and their interactions are known down to the lowest level.

In principle, using robot models rather than animal models may also be preferable because of ethical concerns. At this time, we feel that this is of limited importance because the current level of robotics does not allow detailed modelling of behaviors that are only found in animals of higher phyla, for which strong ethical considerations come into play.

A one-way street – so far. Despite these benefits, however, the flow of information between biology and robotics is at present almost entirely one-directional. While machine builders receive inspiration from biology, examples of significant discoveries in biological systems that were inspired by building robots are, at best, rare. Webb lists some examples in her target article, but they are few and far between. It is not clear whether there are yet any cases in which robot models lead to nontrivial, successful predictions that have been actually confirmed in animals. This is in marked contrast to other modeling techniques, notable especially in computational neuroscience, where computer modelling has become a respected technique among biologists; the surest sign of this being that many experimental groups routinely develop computational simulations themselves.

Why has robotics not been similarly successful? One reason is because the field is still relatively new and small. Biorobotics, in the sense of robots being used to provide insight into biology, arguably started about fifteen years ago with a paper by Brooks (1985), several decades after computational models were first introduced. Furthermore, the number of active researchers in the field is still very small. This is not counting the large number of those who may have completed one robotics project and then reverted to more classical methods. We have encountered a large number of cases of model recidivism, in which a computational model was implemented in hardware but, in the further course of the project, the hardware implementation was abandoned in favor of future development of the computational model. Presumably, it was found that pursuing the hardware implementation is more difficult, expensive, and time-consuming than the implementation of a simulation. This brings us to the second issue, the difficulty of the approach.

Constructing robots is a difficult, expensive process that takes a long time from original design to finished prototype. Moreover, materials are non-standard, and at present, essentially every model has to be developed from scratch. Without doubt, the field would make much faster progress if a robotic equivalent of the PC existed – a low-cost, universally available, and standardized plat-

form allowing rapid prototyping and seamless collaboration of large groups of researchers. A few candidate systems exist, both in hardware (e.g., Lego Mindstorms™ [http://www.legomindstorms.com], the K-team robots [http://www.k-team.com], Tilden's bugs [Haslacher & Tilden 1995]), and in software (e.g., IQR [Verschure & Voeglin 1998]) but at this time it is not clear whether any of these, or any others, will be able to play the role for biorobotics that the PC played for computational modelling.

The way out. While all of this may sound pessimistic, we remain hopeful about the future. First, we believe that systems-level approaches will increase in importance. Second, although we are still far away from the situation in computation where a nearly universal hardware infrastructure is cheaply and readily available, prices of robotic equipment have come down by several orders of magnitude in less than a decade, and the trend continues. Developing hardware models of biological systems may never become a method for everyone, but it will play a larger role once tools become available that will make robotic modeling accessible to a larger part of the scientific world.

ACKNOWLEDGMENT

We thank the National Science Foundation and the Gatsby Foundation for the support of the Telluride Workshop on Neuromorphic Engineering where this work was performed.

The conundrum of correlation and causation

Irene M. Pepperberg

The Media Lab, Massachusetts Institute of Technology, Cambridge, MA, 02139. impepper@media.mit.edu <http://www.alexfoundation.org>

Abstract: Biology can inspire robotic simulations of behavior and thus advance robotics, but the validity of drawing conclusions about real behavior from robotic models is questionable. Robotic models, particularly of learning, do not account, for example, for (a) exaptation: co-opting of previously evolved functions for new behavior, (b) learning through observation, (c) complex biological reality, or (d) limits on computational capacity.

Although Webb presents an important review of robotic models, including excellent guidelines for their biological relevance, she admits that "a model that behaves like its target is not necessarily an explanation of the target's behavior," that is, that correlation is not necessarily causation. Such oft-repeated statements, however, do not dampen her enthusiasm for robotics as a means to understand biological form and function. What her account lacks is additional appreciation of potential problems inherent in using robotics to answer biological questions. She fails to acknowledge exaptation, observational learning, and complexity as biological reality; she underestimates limits on computational capacity.

Exaptation, or the co-opting of previously evolved functions to do new things, can seriously compromise robotic simulation. Evolutionary forces work on existent biology, and thus real-life biological solutions may involve mechanisms less efficient than those used robotically. Hewes (1973), for example, argues that spoken language was derived from gestural forms without major neural restructuring. Data supporting Hewes' hypothesis – and the notion that exaptation of gestural neural substrates for communication may be extremely widespread – are that parallel development of physical and communicative combinatorial acts exists in humans, nonhuman primates, and even Grey parrots (Greenfield, 1991; Pepperberg & Shive 2001). Mechanisms used by a robotic system to model acquisition of spoken language *de novo* might reproduce data, but are unlikely to use circuits derived initially for stacking cups in order to combine labels. Thus, its mechanisms would be removed from, and say little about, those of biological systems.

Observational learning is also widespread in animals (Heyes & Galef 1996). Animals would die before they could reproduce if

they had to learn skills such as predator avoidance or what to eat via the trial-and-error mechanisms that are currently the basis for computer modeling (Pepperberg 2001). Even in the most elegant attempts at imitation simulation, which involve some form of programming by example, the extent to which the computer learns is limited (e.g., Lieberman 2001; Weng et al. 2001). Thus, the current relevance of robotics to forms of learning beyond simple associationist principles, and to real-life systems, is limited.

In a related vein, biology is complex. Advanced learning involves the ability to choose the set of rules, among many learned possibilities, from which the appropriate response can be made, and the creativity to build upon learned information to devise novel solutions to a problem. In contrast, conditioned learning is limited in scope in that it does not allow a robot even the ability to alter behavior quickly based on the immediate past, much less allow immediate flexibility to respond to changing conditions. True, brute force systems such as Big Blue win chess games with stunning success (e.g., Campbell 1996), but such systems cannot learn in a broad manner, that is, cannot integrate new and existing knowledge to solve novel problems, take knowledge acquired in one domain to solve problems in another, or form and manipulate representations to attain concrete goals. The point is not that associative/conditioned learning is irrelevant: It exists, is a basis for learning, can be seen as basic to the programming language of learning . . . ; but it is not the appropriate overall program for learning, because it does not engender generalization, transfer, or insightful behavior. The simple initial association of stimulus and response may be what is first linked in memory in humans, but for humans repeated interactions in the real world both sharpen and broaden the connections (Bloom 2000); what results is a representation. Robots can indeed be programmed so that repeated interactions improve their decision-making ability, and one might even argue that statistically-based similarity coding might constitute a representation. Advanced learning, however, derives from manipulation of representations. What is needed to devise an intelligent learning machine, therefore, is not a more efficient program that takes a stimulus as input and uses various rules to produce an expected response, but one that takes that stimulus and uses creativity, reasoning, and decisions based on context to produce an appropriate, adapted, adaptive behavior. So far, robotics cannot simulate such behavior.

Finally, the computational or robotic capacity used to produce a model might be less than the computational capacity of the living system; we cannot discount real-life mechanisms because simulations cannot reproduce the data. Webb cites Kuwana et al. (1995), who must use the actual antenna of moths on their robot model because available gas sensors are ten thousand times less sensitive than the biological system. Later she comments on the rejection of lobsters' use of instantaneous differences in concentration gradients between their two antennules to do chemotaxis, simply because robotic implementation of this algorithm in the real lobsters' flow-tank failed (Grasso et al. 2000) – that is, she implies that failure could be merely a consequence of the quality of the robotic sensor. I applaud Webb's inferences, but suggest that these problems are more serious than she surmises.

In sum, robotic design can advance from attempts to simulate animal behavior without worrying about simulating exact mechanisms. But using current robotic simulations (which for learning are predominantly based on associationist principles) to answer questions about real-life systems can lead into a trap identical to that of Skinnerian behaviorism, which found many situations these same laws could not explain. Anomalous activities of animals whose natural responses to stimuli could not be reshaped by behavioristic training (e.g., Breland & Breland 1961; see review by Roitblat 1987) required a new paradigm in which animals were seen as multi-level processors of information (Kamil 1984; 1988; Pepperberg 1990). This need was made even clearer by behavioral ecologists, whose data could be explained only by positing mechanisms such as selective attention and long-term memory (e.g., Kamil & Sargent 1981; Pyke et al. 1977; see also Roitblat 1987), which were