Commentary on Barbara Webb (2001) Can robots make good models of biological behaviour?
Behavioral and Brain Sciences (2001) 24(6)

The Methodology of the Artificial

Luc Steels
Sony Computer Science Laboratory - Paris
and VUB AI Laboratory,
Brussels,
1050
Belgium
steels@arti.vub.ac.be
<http://arti.vub.ac.be>

Abstract

Biorobotics research should not only target 'realistic' models of living systems and be judged exclusively from that perspective. It should pay just as much attention to formal models and artificial systems. They allow the examination of assumptions which do not necessarily hold for living systems, but therein lies precisely their value. They generate insight by enabling a comparison between the artificial and the real.

Barbara Webb (Webb, 2001) is to be applauded for her courageous effort to examine the methodological assumptions of those researchers in robotics and artificial intelligence who do not exclusively focus on building practical applications but try to advance scientific knowledge on adaptive behaviour and cognition. Webb focuses on the question in how far robots can be good models of biological behaviour, where "modelling involves the representation or correspondence between a (real) target system and something else" (o.c.). This type of realistic modeling is very common and several robotics researchers reviewed in her paper have stated that their robotic experiments should be interpreted as such. But it is not the only way.

There is a second type of scientific activity which consists in making formal models, as illustrated by work in theoretical biology or theoretical economics. Formal models do not necessarily describe a natural system. But they examine the implications of certain assumptions which then can be used to understand natural systems. The assumptions need not be realistic, on the contrary. Unreality of assumptions is often seen as a virtue: "to be important ... a hypothesis must be descriptively false in its assumptions" says economist Milton Friedman (1953, p.14). Unrealistic assumptions make it possible to investigate boundary conditions, isolate factors, highlight implications which would otherwise go unnoticed, perform demonstrations by *reductio ad absurdum*, etc. Thus Nowak et.al. (2000) make a number of obviously false assumptions in their formal models of the evolution of language, for example that the lexicon and grammar of a language are transmitted genetically. But this does not diminish the strength of their claims about why compositionality may be selectively advantageous. A lot of mathematically oriented research in neural networks and many papers found in conferences on 'simulation of adaptive behavior' or 'artificial life' similarly explore formal models rather than realistic ones.

Building artificial systems is a third and in my view truly alternative type of scientific activity and is one that many biorobotics researchers implicitly practice. It consists in building a machine that has a similar functionality as one performed by a natural system, for example an airplane capable of

artificial flight or an artificial walker capable of walking. An artificial system is on the one hand more 'realistic' than a formal model because it involves building physical systems that undergo the constraints of nature or of the cultural and social environment in which they are put. But the researcher is not restricted to mimicking natural systems. Building artificial systems is therefore similar to formal modeling, it can explore alternative solution paths, use other boundary conditions, adopt pragmatic solutions that are very different from biological implementations. An airplane is not a realistic model of a flying bird. Its wings do not have feathers, it does not flap its wings, it does not run on two legs. Early attempts to build airplanes by modeling biological systems failed miserably. Similarly an artificial walker or an artificial face recognition system do not have to mimic the solutions adopted by human beings, which are hardly understood anyway. What matters is good performance and this can be measured accurately and objectively. Artificial systems require a level of detail which is not necessary in formal modeling. Thus Nowak et.al. (2000) can simply assume that agents learn a grammar with a certain probability without having to specify the learning algorithm itself - whereas an artificial system researching the same questions would need an operational learning scheme that can work on realistic linguistic input and use self-generated meaning anchored in the world through a sensory-motor apparatus (Steels, 1998). Artificial systems are more constrained however, because the solution must work in reality. This gives less freedom in exploration.

Why would we want to use the methodology of the artificial? I see at least three reasons: We may want to replicate the functionality in a cheaper, more reliable, or more robust way. For example we may want to build airplanes to carry people and cargo reliably through the air. We may want to understand how a particular functionality can be achieved at all and thus comprehend a mystery that nature somehow solved. Or, we may want to compare the behaviour and mechanisms of artificial systems to that of natural systems achieving the same functionality. Comparing is not the same as mimicking or modelling the natural system in the sense discussed by Webb (2001). On the contrary, it is a way to gain insight by confronting it with something that is different but still sufficiently similar to make the comparison interesting.

By pretending that the main goal of present-day biorobotics research is to make faithful models of biological systems, we miss out on the opportunities for insight that formal models and artificial systems give us, and we only increase the existing misunderstanding about the methodology of the artificial and hence make it even more difficult for researchers following other approaches to see the point or to learn from our results. There is no doubt that this misunderstanding exists. My current research focuses on language communication, and particularly questions concerning the origins and acquisition of language and meaning (Steels, 1998). I try to get to a point where robots construct jointly artificial languages to communicate about the real world perceived through their sensors. I am often asked whether it is really necessary to build these physical robots, given the effort involved, and whether the same insights cannot be gained by computer simulations. Our papers are occasionally rejected by conferences in linguistics or natural language processing with the argument that the artificial languages constructed by these robotic agents are not natural, as if a sentence constructed for the occasion by a generative linguist or computational language processing without semantics and pragmatics are more natural. In any case, such reactions miss the main points of the methodology, namely (1) artificial systems are developed in the first place to examine the consequences of certain assumptions, just as formal models, (2) they require much more realism than formal models and hence provide much deeper insight, but (3) the goal is not to build realistic replicas of natural systems. Their value for understanding nature lies in providing points of comparison with natural systems. Such an activity is as valuable as making faithful realistic models.

References

- Nowak, M. A., J.B. Plotkin and V.A. Jansen (2000) The evolution of syntactic communication. *Nature* 404 495-498.
- Steels, L. (1998) The origins of syntax in visually grounded robotic agents. *Artificial Intelligence*. 103 (1,2) 133-156.
- Friedman, M. (1953) *Essays in Positive Economics*. University of Chicago Press, Chicago.
- Webb, B. (2001) Can robots make good models of biological behaviour? *Behavioral and brain sciences* (2001) 24 (6).

Acknowledgements

Funding from a governmental GOA project to the VUB AI Lab and CNRS-OHLL project to Sony CSL is gratefully acknowledged. Thanks to Marleen Wynants for editorial comments.