Available online at www.sciencedirect.com

SCIENCE DIRECT®

Cognitive Systems Research 6 (2005) 145–174

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

More things than are dreamt of in your biology: Information-processing in biologically inspired robots ☆

Action editor: Ron Sun

A. Sloman [a,*], R.L. Chrisley [b]

[a] School of Computer Science, University of Birmingham, UK
[b] Centre for Research in Cognitive Science, University of Sussex, UK

Received 22 February 2003; accepted 7 June 2004
Available online 11 September 2004

Abstract

Animals and robots perceiving and acting in a world require an ontology that accommodates entities, processes, states of affairs, etc., in their environment. If the perceived environment includes information-processing systems, the ontology should reflect that. Scientists studying such systems need an ontology that includes the first-order ontology characterising physical phenomena, the second-order ontology characterising perceivers of physical phenomena, and a (recursive) third order ontology characterising perceivers of perceivers, including introspectors. We argue that second- and third-order ontologies refer to contents of *virtual* machines and examine requirements for scientific investigation of combined virtual and physical machines, such as animals and robots. We show how the CogAff architecture schema, combining reactive, deliberative, and meta-management categories, provides a first draft schematic third-order ontology for describing a wide range of natural and artificial agents. Many previously proposed architectures use only a subset of CogAff, including subsumption architectures, contention-scheduling systems, architectures with 'executive functions' and a variety of types of 'Omega' architectures. Adding a multiply-connected, fast-acting 'alarm' mechanism within the CogAff framework accounts for several varieties of emotions. H-CogAff, a special case of CogAff, is postulated as a minimal architecture specification for a human-like system. We illustrate use of the CogAff schema in comparing H-CogAff with Clarion, a well known architecture. One implication is that reliance on concepts tied to observation and experiment can harmfully restrict explanatory theorising, since what an information processor is doing cannot, in general, be determined by using the standard observational techniques of the physical sciences or laboratory experiments. Like theoretical physics, cognitive science needs to be highly speculative to make progress.

☆ Revised version of a paper presented at the WGW'02 workshop on *Biologically inspired robotics: The legacy of W. Grey Walter*, Hewlett–Packard Research Labs, Bristol, August 2002.

* Corresponding author.
E-mail addresses: A.Sloman@cs.bham.ac.uk (A. Sloman), R.L.Chrisley@sussex.ac.uk (R.L. Chrisley).
URLs: http://www.cs.bham.ac.uk/~axs/, http://www.cogs.susx.ac.uk/users/ronc/.

1389-0417/$ - see front matter
doi:10.1016/j.cogsys.2004.06.004

Keywords: Architecture; Biology; Emotion; Evolution; Information-processing; Ontology; Ontological blindness; Robotics; Virtual machines

1. Ontologies and information processing

An ontology used by an organism or robot is the set of objects, properties, processes, etc. that the organism (be it a scientist or a seagull) or robot recognises, thinks in terms of, and refers to in its interactions with the world. This paper discusses some of the components of an ontology required both for an understanding of biological phenomena and for the design of biologically inspired robots. The ontology used by scientists and engineers studying organisms and designing robots will have to include reference to the mechanisms, forms of representation and information-processing architectures of the organisms or robots. Insofar as these natural or artificial agents process information, they will use ontologies. So the ontologies used by scientists and engineers will have to refer to those ontologies. That is, they will have to include meta-ontologies. If we wish to talk about many different organisms or robots (e.g., in discussing evolution, comparing different animals in an ecosystem, or comparing robot designs) our ontology will need to encompass a variety of architectures. At present such comparative studies are hampered by the fact that different authors use different terminology in their ontologies, and produce architecture diagrams using different conventions that make it difficult to make comparisons. In this paper, we present an approach to developing a common framework for describing and comparing animals and robots, by introducing a schematic ontology for some of the high level aspects of a design. We do not claim that this is adequate for all the systems studied in AI, psychology and ethology, but offer it as a first step, to be refined and extended over time.

1.1. Non-physical aspects of organisms and their environments

It is relatively easy to observe the gross physical behaviour of organisms, their physical environment, and to some extent, their internal physical, chemical, physiological mechanisms. But insofar as biological organisms are to a large extent control systems (Wiener, 1961), or more generally *information-processing* systems, finding out what they do as controllers or as information processors is a very different task from observing physical behaviour, whether internal or external (Sloman, 1993, 2003). [1]

That is because the most important components of an information processor may be components of *virtual* machines rather than *physical* machines. Like physical machines, virtual machines do what they do by virtue of the causal interaction of their parts, but such parts are non-physical (by 'non-physical', we do not mean 'not physically realised' or 'made ultimately of non-physical stuff' but merely 'not easily characterised with the vocabulary and methods of the physical sciences'). Compare the notion of a 'propaganda machine'. Entities in virtual machines can include such things as grammars, parsers, decision makers, motive generators, inference engines, knowledge stores, recursive data-structures, rule sets, concepts, plans and emotional states, rather than molecules, transistors or neurones.

An example of a component of a virtual machine in biology is the niche of a species. A niche

[1] Throughout this paper, we use 'information' in the colloquial sense in which information is *about* something rather than in the technical sense of Shannon. That is, like many biologists, software engineers, news reporters, information agencies and social scientists, we use 'information' in the sense in which information can be true or false, or can more or less accurately fit some situation, and in which one item of information can be inconsistent with another, or can be derived from another, or may be more general or more specific than another. None of this implies that the information is expressed or encoded in any particular form, such as sentences or pictures or neural states, or that it is communicated between organisms, as opposed to being acquired or used by one organism. We have no space to rebut the argument in (Rose, 1993) that only computers, not animals or brains, are information processors, and the 'opposite' argument of Maturana and Varela summarised in (Boden, 2000) according to which *only* humans process information, namely when they communicate via external messages.