Behaviorally Coupled Emergent Representation

Joseph Lewis
San Diego State University
5500 Campanile Dr. MC7720
San Diego, CA 92182
1-619-594-2014
lewis@cs.sdsu.edu

Jamie Lawson
Lockheed-Martin Orincon Defense
4770 Eastgate Mall
San Diego, CA 92121
1-858-455-5530
jamie.lawson@lmco.com

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: General

General Terms

Algorithms

Keywords

representation, self-organization, Starcat

ABSTRACT

Traditionally, representation has been perceived as a necessity for producing intelligent behavior. Once the right representation is in place to drive it, behavior unfolds as the system's dynamics interact with what is usually a fixed, structural entity. For many kinds of systems this approach can be successful. However, as a prescription for building increasingly complex adaptive systems, it often fails. An alternative perspective that is under investigation in our Starcat project suggests that representation is not what drives behavior but rather what is left over by the system's dynamics after concepts have been activated and behavior has emerged.

There are numerous examples of patterns emerging from underlying dynamics. In an ant colony, for example, stigmergic behavior arises from the colony's dynamics; but when viewed from outside the system, the pattern reveals the coupling between colony behavior and the environment. We could, from that perspective, consider the pheromone and ant trails as a kind of representation. Particles in cellular automata offer another example that demonstrates how coupling with an external environment (here via the fitness function) draws particular behavior out of the system's dynamics. Again, from an external perspective, these appear to represent information about the environment. Prigogine's dissipative structures describe a similar phenomenon in physical systems far from equilibrium. The familiar Bnard cells could be said to represent a certain level of heat flow through viscous oil. The dynamic nature of these cells further highlights the likelihood that emergent representations are very fluid and likely to change as pressure from the environment changes. Representation in natural systems may well be an emergent phenomenon, a consequence of the system's dynamics, an echo of the coupling between the system's behavior and the pressures from its environment.

There is a family of cognitive architectures, related to Mitchell and Hofstadter's Copycat, which explores these possibilities. The Starcat project attempts to generalize Copycat, bringing several applications under the same design. Starcat is intended to address problems in embodied cognition, where the system interacts with an environment and must produce behavior indefinitely, in the face of changing pressures. It is an architecture for components that produce and consume codelets. The components swim in a virtual sea of different kinds of codelets. The components ignore some codelets and act upon others, while frequently introducing new ones. Some Starcat components couple to the environment, allowing the supply of available codelets to be regulated externally. Each codelet is a short-lived agent that may run and then die. Codelets are by their nature small; and there are many different kinds associated with the system. The primary job of a codelet is to build up or tear down perceptual structures. So codelet activity leaves an echo behind in the form of transient data structures. These data structures "represent" Starcat's perceptions and applications of concepts.

An interesting consequence of Starcat's emergent representation is that the system's myriad micro-behaviors drive the representation rather than, as in traditional systems, knowledge representation driving behavior. Additionally, the coordinated aggregate behavior typical of complex adaptive systems -arising from among the multitudes of interacting local agents- is coupled externally with the environment. In this way, viewed from the outside, the building up and tearing down of microstructures looks like intentional representation.

Knowledge representation in Starcat does not capture concepts, nor does it simply get in the way as Brooks' has asserted. Representation is what is leftover once concepts have emerged. It allows the system to be affected by what it is already doing. Once a behavior is done, the representation can erode because the representation that had built up to support the recently completed behavior is likely to have parts that are irrelevant to the next behavior. New representation soon builds up as part of the next round of behavior, and

GECCO'05, June 25–29, 2005, Washington, DC, USA.

the cycle continues.

The system experiences pressure from outside, and this pressure changes what the system must do to continue to function, even though the specifics of those changes are dictated entirely by its existing internal dynamics. The environment triggers behavior, but it does not specify behavior. This is the notion of autopoeisis. We suspect that at an important relationship exists between autopoeisis, autonomy, and the kind of behaviorally coupled emergent representation that we have been describing. Our various applications and corresponding experiments continue to examine if this is the case.

We have developed a simulation of ant colony behavior that uses a degenerate version of the Starcat architecture. The experiments reveal just how much the microbehaviors of an emergent system can produce macroscopic features that are coupled to the environment. We are about to undertake a new experiment in which the slipnet plays a shaping role for the colony. This may be the lynchpin that ties the external apprehension of representational behavior back into the system. We have an application that will use the perception of patterns in musical input to shape the ongoing structural changes in the slipnet. We also have a set of agents that are each driven by an individual slipnet but which interact with one another in a collective workspace. Finally, we are in the process of reimplementing both Copycat and one of its successors, Madcat, using the Starcat framework.