Making Room for Representation

Seth Bullock

Adaptive Behavior 2003; 11; 279
DOI: 10.1177/1059712303114010

Adaptive **Behavior**

Making Room for Representation

Seth Bullock
School of Computing, University of Leeds

One advantage that a dynamical systems perspective on cognition as situated, embodied action (hereafter abbreviated to "dynamical perspective") is sometimes claimed to have over a more traditional cognition-as-computation approach is its ambivalence with respect to the assumed nature of cognitive innards. Unlike a traditional computationalist approach, a dynamical perspective doesn't automatically assume that a successful explanation of the workings of any cognitive mechanism will reveal how that mechanism relies on the manipulation of quasi-linguistic representations.

But there are at least two different ways to understand this kind of ambivalence. First, it might be understood as a kind of eliminativism. Like a purely physical explanation couched in terms of atomic collisions, a dynamical systems account might simply under-cut high-level cognitive phenomena, and perhaps as a result have no room for notions of internal representation, etc. However, if this were the case, and it was impossible to discover and characterize representational entities through a dynamical systems theory approach, if they simply do not register from this perspective, then there would be little surprise when dynamical approaches failed to generate representation-fuelled explanations. A dynamical systems explanation would always be non-representational even when applied to a traditional good-old-fashioned artificial-intelligence system.

However, the dynamical perspective's ambivalence could be of a second, non-eliminativist kind that might be recast along these lines: while a dynamical systems analysis of a cognitive system *may* lead to the discovery that the mechanisms involved achieve their success without recourse to anything resembling full-blown conceptual representations, it is in principle *possible* that such an analysis might reveal the involvement of an internal "language of thought'' as imagined by traditional cognitive science. In addition to this, proponents of a dynamical perspective would presumably make a second, separate, empirical claim that much natural cognition can be fully explained in the former, non-representational manner.

Beer's (this issue) analysis of his circle-catching agent provides a compelling demonstration of the ability of the dynamical systems perspective to explain the manner in which the coupling between brain, body and world can be shaped to achieve simple discriminatory behavior. Moreover, this explanation is completed without identifying a clear role for traditional cognitive machinery—structure-sensitive processes computing over compositional symbolic representations. However, the absence of anything resembling an internal representation in the explanation of the agent's behavior, while perhaps strengthening the empirical claim that cognition does not have to equal computation, leaves open questions concerning the suitability of the dynamical perspective for revealing the nature of cognitive representations should they ever be encountered. What would representations look like from a dynamical perspective? How would they reveal

Correspondence to: S. Bullock, School of Computing, University of Leeds, UK.
E-mail: seth@comp.leeds.ac.uk;
Tel.: +44-(0)-113-343-5322, *Fax*: +44-(0)-113-343-5468.

(2003), Vol 11(4): 279–280.
[1059-7123(200312) 11:4; 279-280; 041862]

themselves? Would "discovering'' or "identifying'' them require new kinds of dynamical systems technique or approach, or would the analysis presented by Beer suffice?

When considering the kinds of representational phenomena that traditional computationalists might expect to find within a cognitive system, Beer describes the targets as internal agent states that "align in some sufficiently straightforward way" with "objects and relations appearing in our conceptualization of an agent's environment and behavior''. It would appear from this description that if representations are there inside an agent's head, then it will not be too hard to find them. But there is no guarantee that real cognition will meet these expectations. An agent's internal representations may align with objects and relations that play an important role for the agent's internal mechanisms, but do not appear in *our* conceptualization of its environment or behavior. Furthermore, this "alignment" could turn out to be far from straightforward.

Beer also notes the wide middle-ground that stretches between, at one extreme, non-representational internal state that merely correlates with external states of affairs, and, at the other, a traditional notion of representation as symbolic, quasi-linguistic, centralized, compositional, etc. Beer appears to be slightly exasperated by this profusion of representational kinds and to suspect that this activity might simply disguise a fundamental weakness with the entire notion of internal representation. However, if a dynamical systems perspective is truly neutral with respect to the nature of cognitive innards, then it might be profitably used to distinguish between this multitude of different notions of representation and thereby settle some of these issues. Just how useful are these representational notions in understanding natural and artificial cognition?

The thought here is that if Beer's explanatory methodology is unbiased, it could (perhaps rarely) reveal a representational explanation of some particular cognitive behavior. Examining the character of this explanation could lend support to, or perhaps undermine, one or more of the different kinds of representational explanation listed by Beer. Rather than merely taking the recent proliferation of different kinds of representation theory to be the death rattle of representational explanation, the dynamical approach should take seriously the prospect of new and better kinds of representational story. In order for the dynamical approach to do justice to the potential for representational explanation, much hinges on demonstrating that these kinds of story fall within the ambit of dynamical explanation both in principle and in practice.

So, to summarize, when a dynamical systems analysis fails to uncover a role for representation in the workings of a cognitive mechanism, it will be important to understand whether this is because such a representational account is hard (or perhaps impossible) to formulate from a dynamical perspective, or because the particular explanatory methods employed are perhaps biased against such accounts, or alternatively, because there is no need for such a representational account, or such an account does not exist. The challenge for those adopting the dynamical perspective is not simply to provide a representation-free account of cognitive behavior, but to provide one that explains why the behavior appears meaningful, why it is cognition rather than just motion. Whether or not one actually believes that much cognitive behavior relies on the manipulation of internal quasi-linguistic representations, part of the challenge for any cognitive science is to explain, rather than explain away, how cognition comes to appear to be this way.