

But will it scale up? Not without representations.

A commentary on *The dynamics of active categorical perception in an evolved model agent* by R. Beer

Shimon Edelman

Department of Psychology

232 Uris Hall, Cornell University

Ithaca, NY 14853-7601, USA

<http://kybele.psych.cornell.edu/~edelman>

November 28, 2003

Beer's paper devotes much energy to buttressing the walls of Castle Dynamic and dredging its moat in the face of what some of its dwellers perceive as a besieging army chanting "no cognition without representation". The divide is real, as attested by the contrast between titles such as "Intelligence without representation" (Brooks, 1991) and "In defense of representation" (Markman and Dietrich, 2000), to pick just one example from each side. It is, however, not too late for people from both sides of the moat to meet on the drawbridge and see if all that energy can be put to a better use. The parley can be organized around an attempt to identify those attributes of representations that both sides may find useful. In my view, one such attribute is the *capacity for hierarchical abstraction*, an aspect of the representational approach not actually mentioned by Beer.

The capacity for abstraction is gained automatically by those theoretical frameworks that opt for the explanatory benefits of representation in understanding cognition. Moreover, without appropriately structured mediating states, a cognitive system would be incapable of dealing with complex reality. Thus, hierarchical abstraction is useful not just for the cognitive scientists in their attempts to scale up their understanding of cognition: it is also indispensable for cognitive systems that aspire to scale up their understanding of the world.

In his discussion of representations, Beer remarks that “the importance of the states that the system passes through lies not so much in any content that they may be assigned, but rather in their sensitivity to subsequent inputs and the future behavior that they make possible.” In response, I contend that it is in general not possible to characterize “future behavior” without first meeting the challenge of understanding “content” (of states or of trajectories through the state space — it does not matter), unless one is dealing with a toy situation, which is why the issue of scaling up is crucial here.

The realization that complex cognitive systems cannot be understood without resort to a hierarchical abstraction of details has been articulated by Marr, Poggio and their collaborators in the mid-1970s (Marr and Poggio, 1977; Marr, 1982). While Marr’s three distinct levels of understanding — computational, algorithmic and implementational — need not (and probably cannot) be as independent as he envisaged (Edelman, 1999), the framework remains valid because of the central role it assigns to representations, which span the levels of understanding, and afford hierarchical abstraction.

The basic requirement from representations is, as Marr pointed out, to make explicit certain properties of the world, and not merely to play a “direct causal role in the cognitive machinery” (Beer, section 9.3). In the context of contemporary theories of computational cognition, “making explicit” need not mean one semantic label per symbol (a stance that Beer describes as “common”). A more sophisticated and more appropriate form of the basic requirement is that the representations be informative enough to allow the subsequent computations to be carried out by simpler means (thus forestalling an infinite regress and the need for a homunculus).

A beautiful example of this consideration at work in the context of a dynamical systems theory can be found in the ongoing project of Maass, Natschläger and Markram (Maass et al., 2003), who describe a neurally inspired two-tiered approach to cognitive information processing. The first tier consists of an ensemble of neural circuits whose rich and complex dynamics allows it to represent arbitrarily complicated functions of the stimuli. The second tier contains simple decoders that tune in to various aspects of the intermediate representations and map these onto the required outputs. Hierarchical abstraction is manifested in this approach in the encapsulation of the dynamics of the intermediate representations, which effectively become symbols with extremely complex but, mercifully, hidden innards. The growing influence of the related idea of *perceptual symbols* (Barsalou, 1999), which can support learning and generalization through dimensionality reduction and function approximation (Edelman, 1999), undermines Beer’s claim that “. . . the exemplar for computational theories of cognition is still quasilinguistic symbols syntactically manipulated. . .”.

Beer’s anti-representation stance seems to be unwarranted not only because of the meta-theoretical con-

siderations outlined above, but also in the light of his own example of a system evolved to categorize simple shapes. First, the analytical methods he marshals are barely up to the task even in the toy setting of his choice. In this he is in good company: mathematical tools suitable for analyzing complex dynamics in hierarchical, functional terms simply do not exist at present (Bullock, 2003). Second, and perhaps more importantly, the target of the analysis — the evolved solution to the toy task — seems to be hardly worth the effort. In particular, in the example analyzed in section 5.3, the agent trying to deal with a circular shape is described as “completely uncommitted during the initial foveation . . . then largely committed to catching during the first scan (although there are times when this commitment wavers) . . . then uncommitted during subsequent scans, but committed to catching at each reversal . . . The pattern is less clear-cut for diamonds.” To me, these meanderings, described later in section 5.3 rather charitably as “a complex, temporally-extended process”, are an indication of a deep-seated incompetence on the part of the agent. Would you feel safe riding in a car with a vacillating driver whose commitment to crossing a junction goes through a similar sequence of ups and downs?

All this suggests that drawing wide-reaching conclusions about the nature of, and the need for, representations on the basis of a system encouraged not to have any (by being confronted with but a single toy task) merely breeds doubts concerning the ability of the anti-representation theories to scale up. I hasten to stress that dynamical approaches as such are not inherently anti-representation, nor need they be confined to toy examples. Indeed, biologically detailed simulations of dynamical neuronal circuits proved useful in attacking various problems of considerable complexity (Maass et al., 2003). Interestingly, whereas Beer’s approach has a network specially evolved for the particular task at hand, Maass et al. describe an easily trainable, universal architecture, derived from a theoretical framework that not only does not eschew representations, but uses them to explain its own versatility (after all, the richness of internal representations is what allows transfer of expertise between tasks (Intrator and Edelman, 1997; Edelman and Intrator, 2002)).

In conclusion, it looks like dynamical systems (and probably the other components of Beer’s theoretical platform as well, embodiment and situatedness, on which I have not commented) can have a place of honor in computational cognitive science, but not at the expense of the idea of representation. We should remember that there is no understanding of cognition without hierarchical abstraction — simply because there would be no cognition without hierarchical abstraction.

*Acknowledgment.* Thanks to Michael Spivey for his comments on a draft of this paper.

## References

- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22:577–660.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47:139–160.
- Bullock, T. H. (2003). Have brain dynamics evolved? should we look for unique dynamics in the sapient species? *Neural Computation*, 17:2013–2027.
- Edelman, S. (1999). *Representation and recognition in vision*. MIT Press, Cambridge, MA.
- Edelman, S. and Intrator, N. (2002). Models of perceptual learning. In Fahle, M. and Poggio, T., editors, *Perceptual learning*. MIT Press.
- Intrator, N. and Edelman, S. (1997). Learning low dimensional representations of visual objects with extensive use of prior knowledge. *Network*, 8:259–281.
- Maass, W., Natschläger, T., and Markram, H. (2003). Computational models for generic cortical microcircuits. In Feng, J., editor, *Computational Neuroscience: A Comprehensive Approach*. CRC-Press, Boca Raton, FL. to appear.
- Markman, A. B. and Dietrich, E. (2000). In defense of representation. *Cognitive Psychology*, 40:138–171.
- Marr, D. (1982). *Vision*. W. H. Freeman, San Francisco, CA.
- Marr, D. and Poggio, T. (1977). From understanding computation to understanding neural circuitry. *Neurosciences Res. Prog. Bull.*, 15:470–488.