

Mind/body problems? Turn to Beer

Manuel de Pinedo[†] and Jason Noble[‡]

[†]Department of Philosophy, University of Granada, Spain;
Programma “Ramón y Cajal”, MCyT, pinedo@ugr.es.

[‡]School of Computing, University of Leeds, UK, jasonn@comp.leeds.ac.uk.

Beer’s rich and rigorous attempt at dismantling old and deeply entrenched prejudices in cognitive science through the painstaking analysis of an evolved agent is more than welcome—especially so after a decade or more of not-quite-delivered promises from the dynamical systems approach to cognition. We are broadly in sympathy with the perspective Beer espouses, but we want to invite the author to take one more step and join us in seeing his work as leading to consequences that are even more radical than those reached in his paper.

A central part of Beer’s paper is his rejection of representationalism; that is, he rejects the idea that cognition consists of the manipulation of symbolic internal representations. It is important to note that representationalism is just one aspect of a bigger and older idea, namely internalism, the belief that cognition is something that goes on inside the head. Our thesis in brief is that Beer might as well go all the way and reject internalism in all its guises.

Beer seems to take for granted the following dilemma: either higher-level, agent-based explanations (those that use concepts such as cognition, knowledge, meaning, etc.) should be reducible to explanations in terms of internal, representational mechanisms, or such higher level descriptions can be dispensed with in favour of dynamical explanations of the coupled brain-body-environment system. Clearly Beer makes a compelling case for the

second of these two options: after close examination of the circle-catching, diamond-avoiding agent, we find no circle or diamond detectors and nothing that resembles a representation of a circle or a diamond. Appeals to internal representations are not needed to explain the agent's behaviour, whereas the right set of dynamical equations allows us to understand what is going on.

Beer's work is an extended thought experiment with powerful anti-representationalist consequences. Furthermore, Beer is absolutely right to insist that any last-ditch attempt to extend the talk of representations so as to encompass the environment can only lead to confusion. If there are representations, these must be internal, no matter how intricate their causal connection to features of the environment. (Circles and diamonds out there in the world cannot sensibly be described as representing themselves.) It follows then that representations cannot explain perception and action, because, under Beer's view, these are properties of the coupled agent/environment system and not properties of the internal machinery of the agent.

So, we see that representationalism takes something of a battering at Beer's hands. This is all well and good, but we should ask ourselves why representationalist views held any appeal in the first place. The attraction of the position comes about because it matches the Cartesian intuition that thoughts, cognition, and knowledge are things located in the head. Putting it another way, representations are theoretical devices to give flesh to the intuitive appeal of the idea that thought happens inside the agent—in its head, soul, brain, or wherever. Intuitive internalism gives rise to representationalism and not the other way around. Therefore representationalism is best dealt with at its source: that is, by challenging the notion that thought is internal.

Returning to Beer's argument: it seems that he believes that because he has rejected representationalism, he is left with an eliminativist position with respect to concepts like knowledge and meaning. In other words, there is no

place for such concepts in cognitive science, and the only proper explanations will be dynamical systems accounts of cognitive agents coupled with their environments. However, this would only follow if the representationalist perspective was the only way to make sense of knowledge and meaning. We would like to offer an externalist perspective on these concepts which does away with representationalism without disposing of all agent-level talk. (By “externalist” we mean simply the opposite of internalism; an acceptance of the idea that cognition is spread out across the agent and the world.) The dilemma identified by Beer is a false one, and it is possible to save the baby of agent-level explanation while throwing out the representational bathwater. Cognitive science does not need to give up on talking about agent-level concepts such as knowledge and meaning: we just need to recognize that knowledge and meaning, as much as perception and action, are features of the coupled system and not something internal.

Why do we feel that agent-level description is worth saving? Why are we convinced that the impressive analytic tools of dynamical systems theory are not the only tools needed by the cognitive scientist? We refer the reader to the deceptively obvious fact that Beer needs to describe his agent as a circle catcher and a diamond avoider. Indeed, these are the propensities that his agent was selected for over many generations of evolution. This description is admittedly simple, but it is agent-level talk, and clearly of a different explanatory level than a description of the agent/environment system in terms of differential equations. As a quick thought experiment of our own, we ask whether anyone could possibly make sense of the behaviour of the agent given only the dynamical systems description so carefully developed in the paper, and not the brief but enormously helpful agent-level description. Looking only at the lower level, it would be extremely difficult and perhaps impossible to see that all of this complexity was in the service of circle catching and

diamond avoidance.

A parallel observation applies also to the work described in Nepomnyashchikh and Podgornyj's paper. They observe surprisingly fruitful behaviour from systems of coupled oscillators that have received only minimal design effort. However, the behavioural output can only be made sense of in the light of agent-level descriptions such as wandering, searching, and sampling. If systems as simple as those constructed by Beer and by Nepomnyashchikh and Podgornyj require on the one hand agent-level explanations and on the other hand lower-level description in terms of dynamical systems, then clearly cognitive science's more ambitious targets such as human beings and other animals will also require both levels of description.

Thus Beer should not be concerned that endorsing an agent-level description will commit him to the follies of old-fashioned representational cognitive science. One can say that the agent catches circles and avoids diamonds without conceiving of cognition as a series of rule-governed operations over internal symbols. We encourage Beer to go all the way with the externalism he exhibits in his analyses of perception and decision-making. On the view we are urging, mind is not internal, it is all over the place; indeed, "mind" is just a very abstract way of describing the agent/environment interaction. The debate over representationalism loses all urgency once the image of the mind as a place of internal knowings has been dispelled.