




Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Book review

Rolf Pfeifer, Josh Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence*, A Bradford Book, MIT Press, ISBN 978-0-262-16239-5, 2007, 394 pp.

Embodied cognition has been a vibrant and growing research field for around 20 years now, and has produced its share of reviews, overviews and critical assessments (e.g. [5,6,11,2,12,8–10] just to name a few). But despite all of this productive summation, the single best resource for understanding the fundamental motivation and perspective of the field—the one work I always recommend to interested students—has remained Andy Clark's 1996 text *Being There* [3]. Although that speaks well of Clark's book, it is a somewhat unfortunate situation for the field, as a great deal has happened since it was published; indeed, so much—Google scholar lists almost 18,000 articles on embodied cognition appearing since 1996—that it is probably no longer possible to provide a comprehensive overview.

Wisely, Pfeifer and Bongard do not attempt any such global survey of the field, the generality of their title notwithstanding, but rather do *this* community a service by bringing us up to date on what has happened in embodied AI, with a particular emphasis on robotics. They adopt this focus not just for practical, but also for theoretical reasons. For them, building robots is a way—indeed not just a way, but perhaps the best way—of studying intelligence more generally. Why is that? In large part because building robots is an efficient and effective means of focusing attention on the (by now familiar) mantras of embodied cognition: brains are first and foremost control systems for embodied agents, and their most important job is to help such agents flourish. Whether flourishing means following sugar gradients or pheromone trails, climbing fish ladders or corporate ones, success is a matter of flexibly solving the particular problems of survival presented by an agent's environment. Although the ability to manipulate abstract symbols surely has a place in some of the embodied brains that evolved to successfully navigate their environments, it probably isn't a central one, and it certainly isn't the foundation on which anyone would engineer a system capable of solving the problem of survival. At least (and here I am betting that a high percentage of readers are having this very thought) not until that environment starts to have symbols to manipulate, such as can impact flourishing.

If you *were* thinking a thought like that, then you've already accepted most of what Pfeifer and Bongard need in order to make their argument. For whenever in (pre-)history symbols *did* emerge, it is certain that there was a lot of evolving before then, and so even the AI scientist primarily interested in logic, language and mathematics would do well to reflect on what is brought to the table in any symbol fest. What resources might have been exapted, redeployed, recycled or otherwise adapted to these new purposes?¹ And how might that influence what it is we are doing when we seem only to be thinking? Because what intelligence ultimately *is* is the ability to solve problems, understanding it means first knowing what the problems are, in what context they appear, and what other resources (be they bodily or environmental) might be recruited to help solve them. This focus gives their book an interestingly different perspective from most other work in the field. In essence, they ask you to consider: if you were building an agent for a particular environment and form of life, which solutions would best take advantage of environmental resources and physical structure? Would you want a body shaped like this, or like that, and how would each option impact (simplify or complicate) the control system you'd need for the tasks it will perform? This approach means that Pfeifer and Bongard are less focused on the question of how embodiment shapes cognitive processes, the act and content of thinking, and more on how the shape, characteristics and environmental niche of embodied agents impact the design of their cognitive systems. Indeed, they probably should have called their book: “How having *this* body in *that* environment presenting *these* problems determined our mind's design.” It's clunkier than their title, to be sure, but better matched to their content.

As a treatise on the design of cognitive systems, the book is pretty useful. The core, and easily the best part of the book is a long chapter detailing and illustrating eight agent design principles. Each is meant to have a biological basis, and the idea is that if (and only if) we are faithful to these principles when we design intelligent agents, we will begin to understand the true nature of intelligence.

¹ On these issues see, e.g., [1,4].

Agent design principle 1: “Designing an intelligent agent involves the following constituents: (1) definition of the ecological niche, (2) definition of the desired behaviors and tasks, and (3) design of the agent.” (p. 100)

Agent design principle 2: “The complete agent principle states that when designing agents we must think about the complete agent behaving in the real world.” (p. 104)

Agent design principle 3: “The principle of cheap design states that if agents are built to exploit the properties of the ecological niche and the characteristics of the interaction with the environment, their design and construction will be much easier, or ‘cheaper’.” (p. 107)

Agent design principle 4: “The redundancy principle states that intelligent agents must be designed in such a way that (a) their different subsystems function on the basis of different physical processes and (b) there is partial overlap of functionality between the different subsystems.” (p. 113)

Agent design principle 5: “The principle of sensory-motor coordination states that through sensory-motor coordination structured sensory stimulation is induced.” (p. 117)

Agent design principle 6: “The principle of ecological balance has two parts. The first states that given a certain task environment, there has to be a match between the complexities of the agent’s sensory, motor, and neural systems. The second aspect is closely related to the first; it states that there is a certain balance or task distribution between morphology, materials, control and environment.” (p. 123)

Agent design principle 7: “The principle of parallel, loosely coupled processes states that intelligence is emergent from a large number of parallel processes that are often coordinated through embodiment, in particular via the embodied interaction with the environment.” (p. 134)

Agent design principle 8: “The value principle states that intelligent agents are equipped with a value system which constitutes a basic set of assumptions about what is good for the agent.” (p. 137)

Most of these principles are probably clear enough on their own, but a few illustrations will help make their point. Consider, for instance, the problem of controlling the limbs in locomotion. Coherent stepping patterns require precise coordination of the relative timing of limb motions, suggesting the need for something to distribute limb-state information and manage the coordination. According to Pfeifer and Bongard, if we think about this problem from the perspective of principles 3 and 7, we will quickly realize that the body and environment already provide the infrastructure needed to solve this problem, without any centralized control system or communication network. For every time a limb is lifted from the ground, the weight carried by the other(s) changes; similarly, the joint angles of every limb change whenever a limb pushes the body forward. Thus, coordinated walking behavior is possible even when each limb only reacts to local signals, because thanks to the passive transmission provided by body and environment, the local signal also carries global information. Similar considerations can help explain the normal behavior of commissurotomy (split-brain) patients in everyday situations [7]. The globally coherent behavior of a patient with no direct communication links between the two halves of her brain can seem mysterious until we reflect on the fact that the halves are virtually unified by occupying the same body in the same environment, that the behavior of the agent changes the environment and sensory-motor situation in characteristic ways, and that therefore the two halves of the brain are in fact in constant communication, even after losing one important channel.

The notion that acting induces characteristic changes in sensory stimulation is also central to principle 5. For Pfeifer and Bongard, it was a mistake to try to tackle the problem of computer vision by writing software to interpret the signals from stationary cameras (see also principles 1 and 2). Indeed, this apparent “simplification” of the problem, by ignoring the issues of action and locomotion, in fact makes it even harder. In biological systems perception and action are naturally coordinated—that is, every action provides direct (e.g. proprioceptive) and indirect (e.g. visual, haptic) sensory feedback. When I move my arm around the surface of the desk, I don’t just feel my limb move, I see it; and I don’t just see my limb, but I see objects react to this motion (I knock over my drink), and feel their touch (first hard, then cold and wet). Having multiple, intercorrelated sensory streams makes it that much easier to extract signal from noise, imbuing each channel with relevant content and meaning (see principle 8).

In some sense all of their principles really flow from the second, and from the considerations that start to seem central when building complete agents in specific environments. And that, perhaps, is the central value of this book—a reminder that the original vision of AI was not to build special purpose algorithms for solving narrowly circumscribed formal problems, but to create fully, generally intelligent systems. That problem is of course intractable as it stands, and there are two ways to simplify it: analyze the problem and decompose it into tractable parts, in the hope that the parts we build can be assembled into the whole we seek; or identify the most complex whole system we know how to build, in the hope that we can build successively more capable generations of such systems that bring us closer to our ultimate goal. We have spent significant time working on the first option; perhaps it is time to shift more effort to the second. Here we are offered some hope, and some evidence, that such a research program will meet with significant success.

References

- [1] M. Anderson, Circuit sharing and the implementation of intelligent systems, *Connection Science* 20 (4) (2008) 239–251.
- [2] M. Anderson, Embodied cognition: A field guide, *Artificial Intelligence* 149 (1) (2003) 91–130.
- [3] A. Clark, *Being There: Putting Brain, Body and World Together Again*, MIT Press, Cambridge, MA, 1996.
- [4] S. Dehaene, *Reading in the Brain*, Penguin Viking, New York, 2009.

- [5] T. Froese, T. Ziemke, Enactive artificial intelligence: Investigating the systemic organization of life and mind, *Artificial Intelligence* 173 (3–4) (2009) 466–500.
- [6] S. Gallagher, *How the Body Shapes the Mind*, Oxford University Press, Oxford, 2005.
- [7] B. Kolb, I. Wishaw, *Fundamentals of Human Neuropsychology*, Worth Publishers, New York, 2008.
- [8] G. Lakoff, M. Johnson, *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*, Basic Books, New York, 1999.
- [9] R. Pfeifer, C. Scheier, *Understanding Intelligence*, MIT Press, Cambridge, MA, 1999.
- [10] F. Varela, E. Thompson, E. Rosch, *The Embodied Mind: Cognitive Science and Human Experience*, MIT Press, Cambridge, MA, 1991.
- [11] M. Wheeler, *Reconstructing the Cognitive World: The Next Step*, MIT Press, Cambridge, MA, 2005.
- [12] M. Wilson, Six views of embodied cognition, *Psychonomic Bulletin and Review* 9 (4) (2002) 625–636.

Michael L. Anderson
Department of Psychology, Franklin & Marshall College, Lancaster, PA 17604, United States
Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, United States
E-mail address: michael.anderson@fandm.edu

4 August 2009
Available online 17 November 2009