Review

Artificial cognitive systems : a primer
Vernon D., The MIT Press, Cambridge, MA, 2014. 288 pp. Type: Book

Date Reviewed: 06/30/15

Since the earliest days of computers, researchers have sought to use them to emulate human cognition. Some early demonstrations of such "artificial intelligence" (playing checkers, diagnosing diseases) were limited in scope, but the grand goal has always been to encompass the whole range of cognition. The proliferation of attempts makes it difficult for a newcomer to the field to grasp the issues involved. This volume offers a helpful synthesis that spans rival approaches and allows both a high-level survey of the issues and more in-depth references for further research.

The first chapter highlights one of the fundamental challenges in such research: just what do we mean by autonomy? Vernon offers a definition with six facets: autonomy, perception, learning, anticipation, action, and adaptation. This definition is broad enough to encompass most recent research, though some of the facets invite discussion. For example, as long as people can disconnect the power from a device, can that device ever be truly autonomous? For that matter, given the fragile nature of human life in the face of natural disasters, are humans really autonomous? If the answer to both questions is "No," perhaps autonomy is not a requirement for cognition.

The second chapter outlines two main paradigms of cognitive systems that have been developed. Cognitivist systems represent and reason about the external world symbolically, and treat perception and action as discrete events, largely decoupled from the dynamics of the environment. Emergent systems, which include connectionist, dynamical, and enactive approaches, eschew elaborate symbolic representations in favor of close coupling with the environment itself (as in Rodney Brooks' motto, "The world is its own best model"). Throughout the history of cognitive computing, these two approaches have often been in competition, but (as the chapter notes) hybrid systems that draw on both approaches are beginning to appear. A major strength of the book is its persistent comparison of the two approaches in the discussion of each facet of cognition.

Chapter 3 discusses features that should be included in a cognitive architecture, and summarizes three examples: the cognitivist system Soar, the emergent system Darwin, and the hybrid system ISAC. A detailed footnote points to 12 others, leading the interested reader to even more extensive online catalogs.

Chapters 4 through 7 roughly follow the facets of cognition proposed in chapter 1, showing how each is realized in various architectures. The facets of autonomy, learning, and anticipation have their own chapters, while the embodiment chapter covers both perception and action.

Chapter 8 discusses knowledge and representation, a topic to which cognitivist and emergent systems take very different approaches. The discussion of the symbol grounding problem in Section 8.4 touches on many difficult issues, and would have been helped by including P. Gärdenfors' revolutionary proposals [1].

Chapter 9, on social cognition, recognizes that many issues in cognition cannot be understood from the standpoint of a single agent, and must consider cognitive activities that span multiple agents. Additional important references in this area are Wegner's work on transactive memory, and the systematic analysis of multi-agent cognitive configurations by Parunak, Brueckner, Fleischer, and Odell [2].

The book includes detailed footnotes to nearly 400 references that guide the reader who wishes to study further. Frustratingly, the bibliography is numbered in the order in which references appear in the text rather than alphabetically by author, making it very difficult to see quickly which work is being discussed.

This volume fills a serious void in the literature on artificial cognition, spanning a wide range of past work and helping the reader assess the rich array of approaches that have been taken to constructing cognitive systems. It will be an essential text for anyone who seeks to advance work in this field.

1)Gärdenfors, P. *Conceptual spaces: the geometry of thought*. MIT Press, Cambridge, MA, 2000.

2)Parunak, H. V. D.; Brueckner, S.; Fleischer, M.; Odell, J. L. *Agent-oriented software engineering IV (LNCS 2935)*. Springer, , 2004.

Reviewer: H. Van Dyke Parunak Review #: CR143563 (1509-0767)