

Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language

Lucja M. Iwańska and Stuart C. Shapiro (editors)

(Wayne State University and State University of New York at Buffalo)

Menlo Park, CA and Cambridge, MA:
AAAI Press and The MIT Press, 2000,
xix+460 pp; paperbound, ISBN
0-262-59021-2, \$40.00

Reviewed by

Robert E. Mercer

University of Western Ontario

The strong connection between knowledge representation and reasoning and natural language is the main theme of this book. The book is divided into two parts. Part I, consisting of five original or updated papers, considers the connection from the viewpoint of knowledge for language. The six original or updated papers in Part II describe using language to gather or enrich knowledge. Most of the papers discuss significant software systems that are capable of dealing with a variety of corpora. A brief overview of the eleven papers will give some indication of the coverage of the book.

The following five papers constitute Part I of the book. Lucja M. Iwańska, in "Natural language is a powerful knowledge representation system: The UNO system," argues for her conjecture that natural language is a powerful representational language which is particularly suitable for representing and using knowledge in not-well-formalized domains. David A. McAllester and Robert Givan, in "Natural language syntax and first order logic," present a polynomial-time (incomplete) inference procedure for first-order logic given in a nonstandard syntactic representation that takes advantage of properties found in natural language syntax. David D. McDonald, in "Issues in the representation of real texts: The design of KRISP," discusses the need for fundamental changes in our representational languages in order to deal with the analysis and synthesis of real texts in real time. Lenhart K. Schubert and Chung Hee Hwang, in "Episodic logic meets Little Red Riding Hood—A comprehensive natural representation for language understanding," describe their logic for representing discourse content and linguistic and world knowledge that readily captures the meaning of the text and enables appropriate inferences to be made. Stuart C. Shapiro, in "SNePS: A logic for natural language understanding and commonsense reasoning," describes a number of features that are incorporated in the latest version of SNePS.

Although the papers in this part vary somewhat in what they espouse, the primary theme is the value of appropriate representation. It is this aspect which is the important contribution of this section of the book. Although the value of representation is obvious in McAllester and Givan's paper, the theme of their paper is somewhat different from that of the other four, which describe working natural language understanding or processing systems. Shapiro's paper is a terse list of some features found in SNePS 2.4, giving some motivation and a few examples. The remaining three (significantly longer) papers carefully motivate their respective positions and show the applicability of their representational languages with numerous examples.

The next three papers form the first subtheme of Part II of the book, a discussion of uniform and nonuniform representation and reasoning frameworks. These three contributions describe significant natural language systems for machine translation, reasoning about space and time, and dialogue processing in a tutoring system, respectively. Bonnie J. Dorr and Clare R. Voss, in "A multi-level approach to interlingual machine translation: Defining the interface between representational languages," argue that a variety of knowledge representation types is advantageous in their machine translation task. Lucja M. Iwańska, in "Uniform natural (language) spatio-temporal logic: Reasoning about absolute and relative space and time," proposes a uniform knowledge representation and reasoning approach for spatio-temporal reasoning, again returning to her conjecture from Chapter 1 that natural language provides the appropriate representation and reasoning mechanism. Susan W. McRoy, Syed S. Ali, and Susan M. Haller, in "Mixed depth representations for dialog processing," use a uniform knowledge-representation framework to reason about the domain knowledge and the discourse in their tutoring system.

The final three papers constitute the second subtheme of Part II. The underlying theme of two of the papers is the use of linguistic knowledge to enrich knowledge representation or knowledge acquisition techniques (the paper by Iwańska, Mata, and Kruger does not use "a priori existing knowledge" [p. 336] to acquire taxonomic knowledge from texts). Sanda M. Harabagiu and Dan I. Moldovan, in "Enriching the WordNet taxonomy with contextual knowledge acquired from text," discuss their contextualization of concepts in WordNet, enabling pragmatic inferences such as Gricean implicatures to be derived. Lucja M. Iwańska, Naveen Mata, and Kellyn Kruger, in "Fully automatic acquisition of taxonomic knowledge from large corpora of texts: Limited-syntax knowledge representation system based on natural language," employ weak, local-context-based methods to extract taxonomic knowledge from text to be represented using the UNO representation language that was introduced in Iwańska's first paper in the book. William J. Rapaport and Karen Ehrlich, in "A computational theory of vocabulary acquisition," use context (surrounding text, grammatical information, and background knowledge) to learn the meaning of new words (or word senses) encountered in text.

The papers in this collection, for the most part, describe major projects, some that span one or two decades (and for some, their roots can be traced further). The comments that I make below should not be taken to reflect on the authors' contributions to this book.

In the preface, the editors state that this book "contains the most recent theoretical and practical computational approaches to representing and utilizing the meaning of natural language." Unfortunately, the book seems to have taken some time to be published. This delay is evident in a number of ways: Most of the papers have no references beyond 1997 and the few that do appear are self-references. Also, a table in the preface noting some professional activities associated with the topic of the book ends with a 1997 entry. The editors' claim is also undermined by the fact that two of the papers are updated versions of papers from *Expert Systems* 1996 and one is an update of a 1992 *Artificial Intelligence* paper. If the 1997 date is not a coincidence, there has been more recent work done. For example, the publication of Blackburn et al. 1998 marks the beginning of the DORIS project at the University of the Saarland (<http://www.coli.uni-sb.de/~bos/doris/>). DRT-oriented and based on classical logics (representation and inference), it would provide a different perspective on this discussion, somewhat akin to the uniform-nonuniform dichotomy in Part II.

In addition to the eleven papers, three appendices are included "to make this book self-contained" (p. xvii). Unfortunately, Appendix A, entitled "Propositional, First-

Order and Higher-Order Logics," gives eight pages each to propositional and first-order logic (which the editors state are inadequate for natural language) and only half a page to the last topic which is much better suited to the task. Also contributing to the lack of self-containedness (but in a minor way) is the lack of cross-references to the papers contained in the book even though references to prior work by contributors are made (especially in Part I).

The editors might have chosen a slightly different set of papers (four of the eleven papers are authored or coauthored by the editors). The DORIS project, mentioned earlier, would be a candidate for Part I.

Because all of the bibliographies have been moved to the end of the book, I would have liked to have seen the bibliographic entries in the index. Having a reverse index of the references is appealing when the bibliography is presented in this manner.

Reference

Blackburn, Patrick, Johan Bos, Michael Kohlhase, and Hans de Nivelle. 1998. Automated theorem proving for natural

language understanding. *CADE-15 Workshop Problem-solving Methodologies with Automated Deduction*.

Robert Mercer is an Associate Professor in the Department of Computer Science at The University of Western Ontario. His interests include nonmonotonic logics and their use in the representation of natural language knowledge and reasoning. Mercer's address is Department of Computer Science, Middlesex College, The University of Western Ontario, London, Ontario, Canada N6A 5B7; e-mail: mercer@csd.uwo.ca.