Commonsense Reasoning

Erik T. Mueller
(IBM Thomas J. Watson Research Center)

Reviewed by
Boris Galitsky
LogLogic, Inc.

In spite of the recent shift of interest of the computational linguistics community to statistical approaches to NLP, the role of commonsense reasoning in semantic and pragmatic processing remains strong. Traditionally, it has been the commonsense reasoning community, and, in particular, the nonmonotonic reasoning community, that have considered linguistic applications of reasoning; however, these researchers do not go all the way to demonstrate how their formalisms can be implemented as algorithms that would enhance the quality of a particular NLP system.

Regrettably, most commercial search systems do not rely on commonsense reasoning, and as a result it is difficult to access a relevant piece of information if the user is not familiar with domain-specific terms. Furthermore, it is not yet possible to find information on the World Wide Web that is related to a query only by common sense. Erik Mueller's book *Commonsense Reasoning* is a substantial step towards making commonsense reasoning formalisms usable for semantic processing, and it can be used as a good manual on how to enable an ontology to reason about events. Such logical AI notions as *fluent, inertia, pre-condition and effect, nonmonotonicity, circumscription*, and others are well explained and illustrated, so that the respective use cases for linguistic information processing become transparent.

The main question that the book addresses is how to automate commonsense reasoning. The focus of the book is event calculus, which is important, in particular, for answering questions about sequences of actions, augmenting semantic rules and domain ontologies by the event-handling components in question answering, annotation, information extraction, and machine translation systems. The use of the described technique can be straightforward for declarative ontologies, especially those that are logic programming–based. However, substantial design effort would be required to implement these techniques in conventional object-oriented programming languages. Hence this book is a good aid to prepare developers of ontologies and semantic rules for the problems that they will certainly experience.

One of the main problems in presenting the formalisms of logical AI is that they first impress the audience as solving rather trivial cases. It is only when a knowledge engineer starts to implement basic reasoning features and runs into problems building intuitively clear links between facts that he or she starts to appreciate the contribution of commonsense logicians. Being defined by a limited number of first-order logic axioms, event calculus enables a semantic system to reason about a wide range of real-world problems. The book demonstrates that the symbolic approach can deliver elegant design with provable properties. At the same time, logical AI would benefit from contributions from linguists who, having read this book, would find graceful interpretations of commonsense formalisms.

One of the important lessons learned by AI research is that some forms of intellectual activity that we take for granted are extremely difficult to implement as software. The current state-of-the-art AI system can hardly reach the complexity of commonsense reasoning of an 18-month-old child. I believe that Mueller's systematic and thorough introduction to the world of the *HoldsAt* predicate will advance the "mental age" of any CL systems that readers of the book design.

Boris Galitsky is currently involved in infrastructure log analytics as senior computer scientist at LogLogic. He co-founded iAskWeb, a company that provides question-answering services involving commonsense reasoning online, and authored a book on that topic (*Natural Language Question Answering: Technique of Semantic Headers*, Advanced Knowledge Intl, Adelaide, Australia, 2003). Galitsky's address is LogLogic, Inc., 3061-B Zanker Rd, San Jose, CA 95134; e-mail: bgalitsky@loglogic.com.