

Book Reviews

Building Large Knowledge-Based Systems

Ken Levine

Building Large Knowledge-Based Systems (Addison-Wesley, Reading, Massachusetts, 1990, 372 pages, \$39.75, ISBN 0-201-51752-3) by Douglas B. Lenat and R. V. Guha is an interim report on the Microelectronic and Computer Technology Corporation (MCC) Cyc project. Cyc is an ambitious 10-year effort whose goal is to overcome the brittleness of contemporary expert systems by capturing the millions of facts and heuristics that MCC researchers consider to be the consensus reality that all intelligent beings share and that leads to common sense. As the authors state in their preface, "There are deep, important issues that must be addressed if we are ever to have a large intelligent knowledge-based program: What ontological categories would make up an adequate set for carving up the universe? How are they related? What are the important things most human beings today know about solid objects? And so on" (p. xvii).

This book does an admirable job of presenting their research. An engineering attitude is taken throughout; that is, practical solutions are featured without much discussion of theoretical underpinnings (for example, What Yale shooting problem?). Highlights include discussions of the Cyc representation language, Cycl, and the authors' approach to what they consider the important representational thorns that must be faced in building an ultra large-scale knowledge base (ULSKB) system (for example, time, space, causality, uncertainty).

Cycl is based on a combination of a frame system for basic factual knowledge (not only for domain objects but for slots and other metaobjects), a constraint language that is used to enhance the expressive power of the frame system (disjunctions, quantified statements, relationships among

slot values), and 24 inference mechanisms. The inference mechanisms can be run in a forward (cached) or backward (on demand) direction, and each has its own special-purpose truth maintenance system. Any contention between conclusions of these mechanisms is handled through override relationships encoded as slots on the frames that represent the specific instances of these mechanisms.

This approach to knowledge representation (slots and constraints as first-class objects) and inferencing (for example, best-first search using a prioritized agenda of tasks) is much in the spirit of Lenat's work on AM and Eurisco. The authors state that what was holding these programs back was the lack of a large knowledge base to use as a substrate for the learning process; thus, the Cyc effort was begun. Throughout the book, they do a good job of explaining their approach to the overall Cyc ontology as well as specific issues such as individuals versus collections and the use of multiple (internally consistent) models to represent and reason with contradictory information.

The discussion of the inference mechanism is perhaps the strongest section of the book; not only do they justify using a myriad of special-purpose, optimized inference methods in addition to general-purpose ones (for example, forward and backward chaining on general if-then rules and unification on horn clause rules), but they also introduce some important new concepts (more general forms of inheritance and classification, plausible guesses for slot values by resolving constraints). One can see uses for some of these inference techniques independent of their use in Cyc. These techniques might be their most important short-term contribution to the state of the art of expert systems. However, this section would have been more useful to the practitioner if the authors had described some of the trade-offs involved in using these various inference mechanisms.

No applications of the foundations

they build are given, although in relation to their solution to the representational thorns, they state that "each has shown itself to be an adequate way of handling the problem. They have withstood testing by personnel both sympathetic and unsympathetic, skilled and unskilled in various disciplines" (p. 198). The authors also don't give details on how Cyc's structure can be used as a basis for machine learning. Chapter 1 discusses coping with novelty by falling back on analogy and the requirement for a large knowledge base to support this, but the discussion ends here. Similar teasing hints are made about their decision to represent parts of Cycl in (and only in) the Cyc knowledge base, with no further discussion of the possible uses of this reification of Cyc's knowledge. In addition, several of the peripheral issues involved (no pun intended) are either not covered at all (for example, plans for effectors and effectors) or are glossed over (for example, the user interface for knowledge entry and for actually using Cyc).

In general, the authors do an admirable job in leading the reader through the major components of Cyc. The style of writing is relaxed, almost conversational; the book would be a quick read for AI students and professionals, although its lack of references to, and discussion of, previous related work will probably keep it from the wider audience it might otherwise deserve. Much of the discussion is thought provoking and should lead to further work in ULSKB theory and construction.

However, the book does leave much to be desired. Because almost no references are made to previous research, the reader has no feeling for the trade-offs involved in defining knowledge representation and reasoning schemes. Two examples: First, the discussion of the representation of temporal knowledge is extremely interesting and well thought out; however, the previous research that it is built on (for example, D. McDermott,

J. Dean, J. Allen) is not mentioned at all. Second, the authors trace their history of first using numeric confidence factors and then switching to symbolic ones, but no mention is made of the Dempster-Shafer theory, hierarchical Bayesian inference, fuzzy logic, or other conventional reasoning schemes. This approach is typical of this book and limits the audience to those who either have the background knowledge to see the trade-offs the authors were dealing with or to those who wouldn't care. In the worst case, it can be misleading about the state of the art. In addition, in several places in the book, terms and symbols are used well before they are defined; the index, although large, does not seem to be well enough organized to help. Nonetheless, I recommend the book as a supplemental text for an advanced undergraduate course because of the wealth of engineering results presented.

In a small section, the authors discuss their attitude toward the neat versus scruffy approaches to AI. Their purported approach is somewhere in the middle of this spectrum: "Neither of the two approaches is exactly right. We agree that the steps involved in both approaches have to be done, but we believe that neither can be ignored if a proper job is to be done on the other." Their true attitude shows throughout the book: For example, they do not attempt to finesse representation and reasoning problems by working in a circumscribed microworld, but there is no attempt to separate the knowledge-level considerations from the symbol-level considerations (for example, many slots have Lisp functions as their values).

This book is a status report made at the halfway point in this project. Thus far, they have encoded 0.1 percent of the knowledge and are far from the knee of the curve (10 to 50 percent), where useful commonsense reasoning will emerge. They had about one million entries as of May 1989 and expected to increase by a factor of four by mid-1990. As far as speed is concerned, they can process about 2500 primitive (level 0) operations per second. A complicated inheritance operation that causes a reasonable amount of rippling will take several seconds to percolate. Until now, all the knowledge has been entered by hand using a frame editor similar to first-generation

expert system shells. "Our schedule was to have enough of the KB built to transition to natural language understanding as the dominant entry mode by 1994 . . . We're still on schedule" (p. 27).

Lenat states that when he took on this project, he thought that there was a 10-percent chance of success. As of May 1989, he thinks this estimation has been raised to 50 percent. However, I believe the real debate is whether commonsense reasoning will emerge from a massive knowledge base or whether the result of the Cyc project will be more evidence that the standard symbolic reasoning approach to AI is limited in its ability to produce AI. Critics of classic AI, such as H. Dreyfus, have shown the fallacy of assuming that knowledge can be considered independent of our being in the world. Perhaps the future lies in combining a symbolic reasoning system such as Cyc with other technologies (for example, connectionism). Such a combination might be more attractive from a practical and philosophic viewpoint.

Some might take issue with the authors' assumption that the ". . . powerful, elegant set of reasoning methods that form a set of first principles that explain creativity, humor, and commonsense reasoning" (p. xvii) don't exist. I am unconvinced because computer science and AI are still in their infancies. Regardless of this fact, some difficult engineering problems will have to be faced in any move toward machine intelligence. The Cyc effort is the AI community's first attempt at such an effort and, as such, deserves to be taken seriously. Overall, the book is a scruffy presentation of the authors' work but an important experiment and an enlightening and entertaining effort.

Ken Levine is founder and president of Lekton Incorporated. His principal responsibility is as head of the design team for the Lekton Expert Configuration System, a shell for constructing interactive, constraint-based configuration systems.

Manufacturing Intelligence

Richard Stokey

Manufacturing Intelligence (Addison-Wesley, Reading, Massachusetts, 1988, 352 pages, \$43.25, ISBN 0-201-13576-0) by Paul Kenneth Wright and David Alan Bourne develops principles for the design of intelligent

machine tools (for example, lathe, milling machine) for batch manufacturing. To this end, the machinist is recognized as a highly skilled craftsman, and much of the book is devoted to a description of the craftsman's necessary skills as well as the knowledge engineering methods for capturing these skills. The two authors (one in computer science, the other in mechanical engineering) combine their extensive research backgrounds in an attempt to realistically approach this subject.

This is a book of great contrasts that reflects the difficulty of its subject matter, which requires a combination of diverse disciplines and approaches. The story starts with the development of general principles for building an intelligent machine tool; in the second half of the book, the focus shifts to the knowledge engineering of the machinist. The characteristics of the chapters are wildly different: Some contain lucid descriptions of principles and concepts, and others are detailed accounts of sensors and machine tools understandable only to the mechanical engineer. Some chapters contain carefully laid out arguments backed by experimentation and research conducted by the authors and their graduate students, but the last chapter contains a relaxed, speculative discussion between the two authors of the future of manufacturing. Another chapter, which proposes a warmed over blackboard architecture consisting of a team of experts, is anemic and gives only sketchy details. Open problems meant to stimulate graduate research projects are offered at the end of each chapter but seem to entail the creation of large research projects well beyond the resources of a poor graduate student who might appreciate a reduction in the size of the projects. Despite these difficulties and challenges, however, the book, which is presented in four parts, flows well; the authors went to great lengths to present their material in an understandable and organized way, using a combination of flowery language and many tables and figures to summarize their main points.

"The Machine Tool Industry" (part 1) sets the stage for the book by noting trends toward small batch production, especially in the aerospace, ship building, and the tool- and die-making industries, where one-of-a-kind parts with complex shapes must be manu-

factured within strict tolerances from difficult-to-machine materials. The authors persuasively argue for the commercial need for an intelligent machine tool in these industries and then give the goals and objectives of their research.

"Building Intelligent Machines" (part 2) is concerned with the subsystems of an intelligent machine tool and their unification. The automated machining of a part requires the cooperative actions of a vision component, a robot hand to place the part, a programmable jig to hold the part, and a machine tool to do the actual cutting. Currently, these different pieces of hardware are made by separate vendors, hence the need for a unifying architecture to handle communications and cooperative problem solving.

An immediate solution is offered for the more short-term problem ("CML: A Meta-Interpreter for Manufacturing," *AI Magazine*, Fall 1986) of translating commands between machines built by multiple vendors. The authors do not leave off here, however; they go on to discuss the fundamental problem (that will not go away by having one integrated system built by a single vendor) of unifying the different representations of the machine tool subfunctions. The goal of identifying a part, for instance, could be done by the coordinated actions of the vision system and the robot hand. To accomplish this task, one must know the equivalent representation for a search goal in vision space and grasping space and be able to translate between the two. For example, the vision system scans a scene acquiring a centered closeup of the object, while the robot hand feels the environment, finally grasping the object. The authors discuss other equivalent representations that occur during the performance of generic tasks. Two chapters in part 2 are devoted to vision and the robot hand, respectively. These chapters are more of a summary of current design rules for building these systems and do not relate to the original goal of designing an intelligent machine tool. The details of other sensory and control modalities are ignored.

In "The Skills of the Craftsman" (part 3), the authors introduce the reader to the true complexity of the machinist's job and what its automation would involve. They carefully construct an argument that heuristic

knowledge is a necessary part of an intelligent machine tool. It is briefly mentioned that some of the requirements for the machinist's sensory and control skills might eventually be eliminated through advancements in material and cutting methods. A discussion here of which skills could be eliminated by such advancements and what fundamental problems remain would have been useful. The authors go on to describe interviews they conducted with machinists using a combination of protocol analysis and videotaping techniques. This information is used as the basis for the automatic setup planner discussed in a separate chapter. Another chapter is devoted to the interesting problem of automating clamps, jigs, and fixtures. Although this problem seems to be mostly a mechanical engineering one, planning is still necessary for the setup and placement of the fixtures.

In "Autonomous Manufacturing" (part 4), both authors make predictions about the factory 30 years from now, one from a computer science perspective, the other from a mechanical engineering perspective. This chapter will be enjoyed by those who appreciate speculation and discussions about long-range research goals and can be safely omitted by those who do not. The glossary helps clarify terms for the reader, and a brief bibliography gives references for further reading.

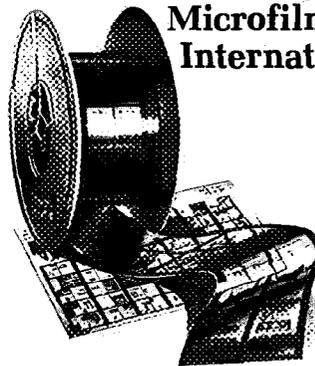
This book has two separate goals: (1) developing general principles for building intelligent machines and (2) automating the machinist's skills. To accomplish the first goal, the authors have woven research in the areas of vision, robotics, manufacturing science, and computer science to present a unified view of the essential subsystems needed for factory automation. Although this undertaking is an ambitious one for any book, the result is an eloquently organized summary of current work. However, I found most of the original contribution of this book to be in the second goal. The authors break new ground by conducting an empirical study of the machinist in action and then presenting prototype systems based on this analysis that plan part setup and automate fixtures and cutting sensors. The chapter on programmable fixtures was particularly interesting because this companion to the machine tool is necessary for machining but is not often considered for

automation. Coming from a computer science background, I found it useful to deepen my knowledge of manufacturing (especially in preparation for part 3) by reading two references supplied in the bibliography.

In their quest to design an intelligent machine tool, the authors take a realistic, empirical approach. They demonstrate that an intelligent machine tool requires a tightly coupled interplay between intelligent software and intelligent hardware. This fact is clarified by distinguishing between the role for verbal knowledge expressible in software (linguistic knowledge) and the construction of flexible, adaptive hardware (non-linguistic knowledge). I recommend this book to AI researchers not only for its valuable insights into an important manufacturing problem but also for its original solutions to this problem.

Richard J. Stokey is a researcher in the Automation Sciences Laboratory at Northrop Research and Technology Center, Palos Verdes Peninsula, CA 90274.

**This publication is
available in microform
from University
Microfilms
International.**



Please send information about these titles:

Name _____

Company/Institution _____

Address _____

City _____

State _____ Zip _____

Phone () _____

Call toll-free 800-521-3044. In Michigan,
Alaska and Hawaii call collect 313-761-4700. Or
mail inquiry to: University Microfilms International,
300 North Zeeb Road, Ann Arbor, MI 48106