The Knowledge Engineering Review

A strategy for near-term success using knowledge-based systems

S.C.LAUFMANN

Pacific Northwest Laboratory, Richland, Washington, USA*

Abstract

Knowledge-based system (KBS) technologies have been applied to a variety of knowledge-related tasks with varying degrees of success. Differentiating among classes of knowledge-related tasks, based on the amounts of problem-solving knowledge and case-specific data involved, can provide valuable insight into why this occurs. Based on this comparison, four classes of problems are described. One class, of data-intensive tasks, includes problem types that are difficult or impossible for humans to perform, yet may be solved in a cost-effective manner using currently accessible KBS technology. The characteristic features of problems in this class are given, together with an example of a successfully fielded knowledge-based system that solves a problem from this class.

1 Introduction

Many successful knowledge-based systems (KBS) have been reported in the research literature. "Classical" systems such as MYCIN (Buchanan and Shortliffe, 1984) and XCON (McDermott, 1982) are frequently referenced. These systems focus on a subset of all knowledge-related tasks, namely those tasks that are concerned primarily with knowledge representation and access issues. Such systems tend to be measured by the quantity of knowledge required to accomplish the task or solve the problem (Buchanan, 1986). This sort of one-dimensional view of KBS technology can obscure issues that may have equally significant impacts on the application of KBS technology in commercial and industrial environments.

To broaden this view, this article considers an additional dimension for evaluating KBS technologies. With a view toward the quantities of case-specific data involved in solving KBS tasks, an emerging class of KBS applications becomes apparent. The four classes of problems identified in this two-dimensional analysis are described in the following section, each with different characteristics and each best-suited to a different KBS approach. Next, the rationale behind this analysis is given, including a view of KBS technology from a pragmatic perspective. One of the four classes is identified as holding particular promise for providing near-term success for moderately-sized commercial and industrial organizations. A candidate task from this class is described, as is the resulting KBS solution, known as the Logger Verification Assistant (LVA) (Laufmann and Crowder, 1987). This system was developed and fielded based on the view given in this article.

2 Classes of problems

An analysis of the amount of problem-solving knowledge versus the amount of case-specific data involved in various types of candidate problems to be solved using KBS technology provides useful insight into problem classification. Figure 1 gives a qualitative view of this comparison. Although no adequate metric is known for quantifying knowledge, for practical purposes the knowledge units

Current address: US WEST Advanced Technologies, 6200 South Quebec Street, Suite 240, Englewood, Colorado, USA
* Operated for the United States Department of Energy by Battelle Memorial Institute.

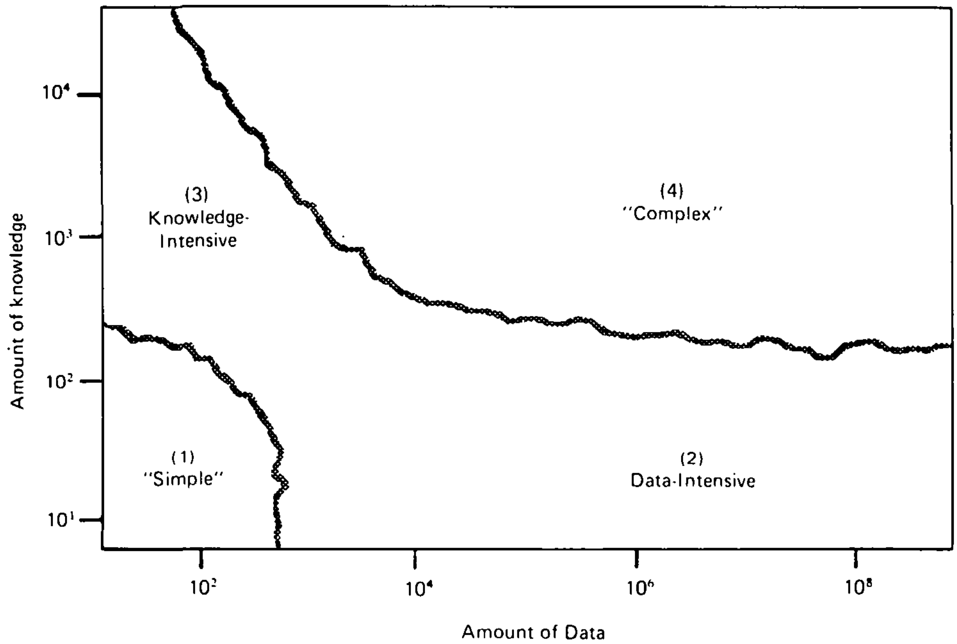

Figures 1 Qualitative comparison of the amounts of problem-specific knowledge *vs.* case-specific data in knowledge-based applications

may be interpreted as the number of rules in a rule-based representation. Data units correspond roughly to the number of case-specific data points that must be accessed by the KBS.

With this view there are four classes of problems, each exhibiting different characteristics and each suited to a different KBS approach. The "simple" problems fall into class 1. These problems involve small amounts of both knowledge and data, and are typically solved with a straightforward rule-based approach. Class 2 includes problems that are characterized by difficult data-handling issues, with less emphasis on knowledge. Systems designed for these problems focus primarily on database solutions, with relatively small knowledge bases. The problems included in class 3 are characterized by difficult knowledge engineering or representation issues, with less emphasis on data issues. Systems designed for these problems tend to involve relatively small case-specific data sets and have simple data-related mechanisms. Most systems reported in the KBS research literature fall into this region. A fourth class encompasses complex problems that entail difficulties related to both knowledge and data issues. These four classes are qualitatively distinct, but lack clear, quantitative boundaries. Thus, problems that lie in the areas between regions may exhibit combinations of the characteristics described below.

The Pacific Northwest Laboratory has developed KBS solutions to various problems in classes 1, 2, and 3. The following characterizations of solutions to these problem classes are based on the results of these implementation efforts, together with research literature (Bachant and McDermott, 1984; Ganascia, 1986; Gaschnig et al., 1983; Hart, 1986; Kaplan, 1986; Kraft, 1984; Smith, 1984; Soloway, Bachant and Jensen, 1987), and the popular literature.

Class 1 systems involve simple, knowledge-related tasks. Such systems have the following general characteristics.

- Low-cost development tools or shells generally provide adequate environments for solving problems in this class.
- The KBS can often be implemented by individuals who lack extensive training.
- Solutions are relatively inexpensive.
- Solutions have low expected value, due to the inherent simplicity of the problem.

Class 2 systems have very large data sets. The characterization of these systems is as follows.

- Mid- to high-cost commercial KBS development tools will usually provide an adequate development environment. Low-cost tools may suffice if they incorporate adequate facilities for accessing external databases.
- Moderately well trained developers can effectively solve these problems.
- Solutions are within the limits of feasibility for moderate-sized organizations.
- Solutions have moderate to high expected value, and may out-perform humans at the data-released tasks.

Class 3 systems are characterized by the high costs of knowledge engineering and knowledge maintenance associated with the leading edge of KBS technology. Perhaps the most mature and best documented class 3 system in industry is XCON, also shown as R1, implemented for Digital Equipment Corporation at Carnegie-Mellon University by John McDermott (McDermott, 1982). Class 3 systems are typically aligned with one or more of the KBS research issues, such as truth maintenance, learning, reasoning strategies, novel representation strategies, and new search techniques. Such systems are characterized as follows.

- Commercial tools are generally inadequate for such problems, so implementations must be based on modified or extended commercial tools, or coded from scratch.
- Highly trained research-oriented individuals, working in groups, have the greatest probability of success.
- Solutions are relatively expensive to develop.
- Solutions have high expected value, but may fail to perform as well as or as fast as human experts in some tasks.

Class 4 problems are complex, because they combine large KB systems and large databases. Solutions to these problems frequently involve research issues from both the KB and the database fields. Class 4 systems are characterized by the following.

- Commercial tools are inadequate.
- Research-oriented developers, most likely in both KBS and database research, are required.
- Solutions are expensive to develop and operate.
- Solutions have high expected value, but implementations may be inefficient.

3 The context for KBS technology

According to Patrick H. Winston, "The primary goal of artificial intelligence is to make machines smarter. The secondary goals of artificial intelligence are to understand what intelligence is (the Nobel laureate purpose) and to make machines more useful (the entrepreneurial purpose)" (Winston, 1984).

Traditional computing without KBS techniques has produced prodigious results, routinely out-performing people in many non-trivial tasks. The primary strengths of "traditional" computing include numerical manipulation, processing speed, consistency of approach, repeatability of results, thoroughness, reliability, and the ability to be copied any number of times. Humans are lacking in these areas, but exhibit proficiency at symbolic manipulation, reasoning, creativity, insight, judgment, common sense, and learning. By merging KBS technology with the traditional computer strengths, it is possible to make "smarter" machines that are capable of out-performing expert humans at some tasks. In so doing, KBS technology extends the applicability of computer systems to new and previously unapproachable classes of problems (Kaplan, 1986). Knowledge-based system technology may therefore be viewed as another step in the ongoing effort to increase the utility of machines.

This has led to an understandable excitement in the business community. However, the promise of KBS technology has not yet been realized fully. Successes to date have typically been class 3 systems. Because many of these systems were developed for KBS research purposes, this approach makes sense. However, industrial and commercial organizations apply different constraints in determining which problems are important and what resources should be dedicated to their solutions. These organizations are looking for ways to achieve success in the near term with fewer, less research-oriented resources.

With this view of knowledge engineering as an incrementally developing technology, the reasonable approach is to concentrate current work on those problems that will yield to currently accessible technology, and to postpone development on more difficult problems until the technology becomes better understood.

4 Toward near-term success

The commercial KBS developer may, therefore, reasonably focus on class 2 problems. Although class 1 systems are relatively simple to implement, experience has shown that they are frequently limited in their usefulness. Class 3 and class 4 systems involve high development and fielding costs, making them generally inconsistent with the goal of near-term success with reasonable costs.

Class 2 tasks involve problems that are data-intensive, requiring relatively small amounts of knowledge to solve a problem that involves relatively large amounts of data. Such problems can be solved without great difficulty given current KBS technology, yet can address useful tasks that may be impractical, too tedious, or overly time-consuming for humans. This class of problems holds great promise for near-term exploitation by industrial and commercial organizations. An example of a class 2 system is given by Utt et al. (1987).

5 A candidate task and resulting system

A typical candidate task from class 2 requires the user to apply some kernel of an intelligent human quality (e.g., heuristic analysis) to a problem requiring the assimilation of an impractically large amount of data. An example is a task requiring an analyst to review 50 pages of numerical data to determine which among a set of process switches is improperly set. A knowledge-based solution to this problem may be reasonably simple to implement, yet the KBS may perform faster and more accurately, consistently, and reliably than the "expert" analyst.

A KBS successfully developed to solve such a problem is the Logger Verification Assistant (LVA), which diagnoses installations of digital data loggers. This system examines the data generated by the data loggers to determine the correctness of the installation. If a fault is detected, the LVA attempts to isolate and identify the specific location and cause of the fault. This system concentrates on the most common, easily identified fault types, leaving the more difficult, unusual types for a human expert. By finding the "easy" faults and presenting only the more challenging cases to the expert, the LVA significantly improves throughput, while reducing the expert's tedium. This enhances the expert's level of expertise, and can allow additional faults to be added to the "easy" list as they become better understood. The LVA performs as a "competent assistant" rather than as an "expert". The knowledge engineer used the domain expert to describe and formulate those more common, simpler heuristic tasks, rather than trying to replicate the highest levels of expertise.

The LVA was developed by a single knowledge engineer in approximately 11 man-months. About 6 man-months were dedicated to problem-solving knowledge issues, and 5 man-months to database issues. The system reviews tens of thousands of numeric data points, but has a relatively small quantity of knowledge. The current rule count is about 200, with approximately 50 additional frames for representing domain-specific conceptual knowledge. The knowledge in this system deals primarily with heuristic diagnostic analysis. Laufmann and Crowder (1987) describe the system in greater detail. This system performs well, correctly diagnosing most occurrences of the basic faults and increasing the expert's throughput by approximately 40%.

Other approaches to solving class 2 problems are examined in Missikoff and Weiderhold (1986).

6 Potential benefits

Focusing KBS attention for the near term on class 2 data-intensive problems can provide businesses and organizations with cost-effective, useful knowledge-based systems that solve significant problems. Experience with the LVA system reveals other benefits that may be common to many class 2 problems.

- The overall solution out-performs humans at some data-intensive tasks.
- The KBS solution may be implemented by existing science and engineering staff, with moderate levels of training.
- Many of the benefits of KBS technology, such as increased throughput, thoroughness, and consistency, can be attained without incurring overwhelming costs, because the artificial intelligence portion of the system remains within the bounds of the more mature, well understood KBS technologies.
- The overall solution integrates traditional programming, directed towards issues such as database access, with a relatively small KBS. When the traditional programming already exists within the organization, development costs are further reduced.

Acknowledgements

Special thanks to Mark Whiting, Mike DeVaney, Erik Pearson, Ron Melton, and Scott Crowder for thought-provoking discussions and suggestions in the preparation of this article.

References

Bachant, J and McDermott, J, 1984. "R1 revisited: four years in the trenches" *The AI Magazine* 5(3) pp. 21–32.
Buchanan, B G and Shortliffe, E H Eds., 1984. *Rule-Based Expert Systems, The MYCIN Experiments of the Stanford Heuristic Programming Project.* Reading, Massachusetts: Addison-Wesley.
Buchanan, B G, 1986. *What's Doable in Knowledge-Based Systems?* Presented at the Fifth National Conference on Artificial Intelligence, August 11–15, 1986, Philadelphia, Pennsylvania.
Ganascia, J-G, 1986. "Using an expert system in merging qualitative and quantitative data analysis" In: *Knowledge-Based Problem Solving*, Kowalik, J Ed. Englewood Cliffs, New Jersey: Prentice-Hall.
Gaschnig, J, Klahr, P, Pople, H, Shortliffe, E, and Terry, A, 1983. "Evaluation of expert systems: issues and case studies" In: *Building Expert Systems*, Hayes-Roth, F, Waterman, D A, and Lenat, D B Eds., pp. 241–280. Reading, Massachusetts: Addison-Wesley.
Hart, P E, 1986. "Artificial intelligence in transition" In: *Knowledge-Based Problem Solving*, Kowalik, J Ed., pp. 296–310. Englewood Cliffs, New Jersey: Prentice-Hall.
Kaplan, S J, 1986. "The industrialization of artificial intelligence" In: *Knowledge-Based Problem Solving*, Kowalik, J Ed., pp. 311–324. Englewood Cliffs, New Jersey: Prentice-Hall.
Kraft, A, 1984. "XCON: an expert configuration system at digital equipment corporation" In: *The AI Business*, Winston, P and Prendergast, K, Eds. Cambridge, Massachusetts: MIT Press.
Laufmann, S C and Crowder III, R S, 1987. "LVA: a knowledge-based system for diagnosing faults in digital data loggers" In: *Proceedings of the Third Conference on Artificial Intelligence Applications*, Silver Spring, Maryland: IEEE Computer Society Press.
McDermott, J, 1982. "R1: a rule-based configurer of computer systems" *Artificial Intelligence* 19(1) pp. 39–88.
Missikoff, M and Wiederhold, G, 1986. "Towards a unified approach for expert and database systems" In: *Expert Database Systems, Proceedings of the First International Workshop*, Larry Kerschberg, Ed. Menlo Park, California: Benjamin/Cummings.
Smith, R G, 1984. "On the development of commercial expert systems" *The AI Magazine* 5(3) pp. 61–73.
Soloway, E, Bachant, J and Jensen, K, 1987. "Assessing the maintainability of XCON-in-RIME: coping with the problems of a VERY large rule-base" In: *Proceedings of the Sixth National Conference on Artificial Intelligence*, Los Altos, California: Morgan Kaufmann.
Winston, P H, 1984. "Perspective" In: *The AI Business*, Winston, P and Prendergast, K, Eds. Cambridge, Massachusetts: MIT Press.
Utt, W K, Anderson J B, Bollag, F J and Nystrom, S D, 1987. "An expert system for data reduction" In *Proceedings of the Second Conference on Artificial Intelligence Applications*, Silver Spring, Maryland: IEEE Computer Society Press.