

Limits to Knowledge Level-B Modeling (and KADS)

Tim Menzies

Department of Software Development, Monash University
Caulfield East, Melbourne, Vic, 3145.; `timm@insect.sd.monash.edu.au`

August 1, 1995

Abstract

Despite the current enthusiasm for knowledge level-B modeling in general and KADS in particular, we find there exists little empirical proof of the utility of this approach. In this paper, we will review the available evidence to suggest that alternative, less abstract, approaches may be better than KADS.

1 Introduction

In the 1970s and early 1980s, several high-profile expert system successes were documented: e.g. MYCIN [49], CASNET [45], PROSPECTOR [6, 14], and XCON [2]. However, despite careful attempts to generalise this work (e.g. [42]), expert systems construction remained a somewhat hit-and-miss process. By the end of the 1980s, it was recognised that our design concepts for knowledge-based systems were incomplete [5].

A new expert system design approach (which has come to dominant the knowledge acquisition (KA) field) is the the search for reusable abstract domain-independent problem-solving strategies. We call this approach \mathcal{KL}_B since it is a variant of Newell's *knowledge level* (KL) modeling approach [28, 30, 29]. KADS is a \mathcal{KL}_B variant used by many commercial expert systems practioners¹. Wielinga *et. al.* note that, as of 1992, KADS has been used in some 40-to 50 KBS projects, 17 of which are described in published papers [46]. Amongst AI-94 participants, interest in KADS was high.

Despite the current enthusiasm for \mathcal{KL}_B and KADS, we find their exists little empirical proof of the utility of this approach. If we review the available evidence, we find that alternative, less abstract, approaches are arguably more productive than KADS. We therefore believe that we should retreat from high-level \mathcal{KL}_B approaches to techniques based on a much lower level of abstraction.

2 A Tutorial on Knowledge-Level Modeling

In Newell's KL approach, intelligence is modeled as a search for appropriate *operators* that convert some *current state* to a *goal state*. Domain-specific knowledge as used to select the operators according to *the principle of rationality*; i.e. an intelligent agent will select an operator which its knowledge tells it will lead the achievement of some of its goals. If implemented, this KL is built on top of a *symbol-level* containing data structures, algorithms, etc. However, to a

¹See <http://www.swi.psy.uva.nl/projects/CommonKADS/Reports.html> for on-line versions of the KADS documentation. See [43] for a detailed text on KADS. See [20] for a good tutorial introduction to KADS.

KL agent, these sub-cognitive symbol-level constructs are the tools used “sub-consciously” as it performs its KL processing [28].

Newell’s subsequent exploration of the KL lead to a general rule-based language called SOAR [38] which was the basis for the problem-space computational model (PSCM) [48]. Programming SOAR using the PSCM involves the consideration of multiple, nested problem spaces. Whenever a “don’t know what to do” state is reached, a new problem space is forked to solve that problem. Newell concluded that the PSCM was the bridge between SOAR and true KL modeling [30, 29].

We distinguish between PSCM (which we term \mathcal{KL}_A) and \mathcal{KL}_B , a KL-modeling variant which groups together a set of authors who argue for basically the same technique; i.e. Clancey’s model construction operators [9], Steels’ components of expertise [40], Chandrasekaran’s task analysis, SPARK/ BURN/ FIREFIGHTER [21] and KADS [46]². The fundamental premise of \mathcal{KL}_B is that a knowledge base should be divided into domain-specific facts and domain-independent problem solving methods. For example, Clancey argues that knowledge engineering should separate heuristics like Figure 1 into domain-specific assertions about the terminology (see Figure 2), re-usable problem solving strategies (see Figure 3) and true domain-specific heuristics (see Figure 4) [9]. Such problem-solving strategies are implicit in \mathcal{KL}_A . The observation that a PSCM system is performing (e.g.) classification is a user-interpretation of a lower-level inference (operator selection over a problem space traversal) [48].

```

if    the infection in meningitis and
      the type of infection in bacterial and
      the patient has undergone surgery and
      the patient has undergone neurosurgery and
      the neurosurgery-time was < 2 months ago and
      the patient received a ventricular-urethral-shunt
then infection = e.coli (.8) or klebsiella (.75)

```

Figure 1: A rule. Domain terms are underlined.

```

subtype(meningitis, bacteriaMeningitis).
subtype(bacteriaMeningitis, eColi).
subtype(bacteriaMeningitis, klebsiella).
subsumes(surgery, neurosurgery).
subsumes(neurosurgery, recentNeurosurgery).
subsumes(recentNeurosurgery, ventricularUrethralShunt).
causalEvidence(bacteriaMeningitis, exposure).
circumstantialEvidence(bacteriaMeningitis, neurosurgery).

```

Figure 2: Domain knowledge from Figure 1.

In the \mathcal{KL}_B view (which we believe is overly-optimistic), knowledge engineering becomes a structured search for an appropriate problem solving strategy. Once a strategy is found or developed, then KA becomes a process of filling in the details required to implement that strategy.

²See the *Related Work* section of [46] for a discussion of the differences in these techniques

<i>Strategy</i>	<i>Description</i>
<i>exploreAndRefine</i>	Explore super-types before sub-types.
<i>findOut</i>	If an hypothesis is subsumed by other findings which are not present in this case then that hypothesis is wrong.
<i>testHypothesis</i>	Test causal connections before mere circumstantial evidence.

Figure 3: Problem solving strategies from Figure 1.

```

if    the patient received a ventricular-urethral-shunt
then  infection = e.coli (.8) or klebsiella (.75)

```

Figure 4: The true heuristic contained in Figure 1.

Libraries of problem solving strategies become a productivity tool for building a wide-variety of expert systems. All known problem solving strategies (e.g. prediction, monitoring, diagnosis, cover & differentiate, propose & revise, planning, design, verification, assessment) are really combinations of a small number (say, ≤ 20) of reusable inference subroutines (e.g. instantiate, generalise, abstract, specify, select, assign-value, compute, compare, match, assemble, decompose, transform). New problem strategies can be quickly built out of these lower-level inference primitives. Knowledge engineers can now separate their analysis work from their design and coding work. Analysis is the process of generating abstract problem solving strategies. Design and coding is a process of implementing these strategies.

\mathcal{KL}_B is not a theory of human cognition. The current public line in the KA community is that knowledge bases are inaccurate approximate surrogate models of reality [12, 17]. Gone are the days of “expertise-transfer” where KA workers believed they were “mining the jewels in the expert’s head” (to use Feigenbaum’s poetic phrase [16]). While some knowledge representation theorists still make occasional claims that their knowledge representation theory has some psychological basis (e.g. portions of the \mathcal{KL}_A research reported in [38, page xxvii]), the official public line is that representations are models/surrogates only [32].

3 Assessment of \mathcal{KL}_B

3.1 Evidence For \mathcal{KL}_B ?

Marques *et. al.* report significantly reduced development times for expert systems using a library of 13 reusable inference subroutines (eliminate, schedule, present, monitor, transform-01, transform-02, compare-01, compare-02, translate-01, translate-02, classify, select, dialog-mgr) in the SPARK/ BURN/ FIREFIGHTER (SBF) environment. In the nine studied applications, development times changed from one to 17 days (using SBF) to 63 to 250 days (without using SBF) [21]. To our knowledge, this is the largest documented evidence of productivity gains in any software approach (be it knowledge based or otherwise). However, in the SBF experiments, the mappings between a graphical language for expressing the specification to the library of inference subroutines was hand-coded. Techniques to assist/automate this mapping process have yet to evolve. The crucial test for SBF is how well this system ports to domains outside of the areas it was initially developed for. We find it significant that the SBF group made no

entry to the Sisyphus-1 or Sisyphus-2 KA project (Sisyphus is an attempt by the international KA community to define reproducible KA experiments [20]). Most other \mathcal{KL}_B groups offered solutions to the 1991, 1992, and 1994 Sisyphus rounds. There was even a \mathcal{KL}_A contribution to the 1994 Sisyphus round [47]. If SBF was a general productivity tool, then it should have been able to quickly build Sisyphus solutions.

Clancey’s classic *Heuristic Classification* paper [7] offered a unified retrospective view on numerous, seemingly different, expert systems. Similar (but smaller) studies (e.g. [1, 20]) suggest that \mathcal{KL}_B can retrospectively clarify historical expert systems design issue. However, two important issues remain unproven: (i) that \mathcal{KL}_B can simplify a *current* developing design; and (ii) that \mathcal{KL}_B is a *comparatively better* design methodology than other approaches.

The separation of problem solving methods from domain assertions and true heuristics (e.g. Figure 1 to Figure 4) can optimise knowledge base processing. Clancey reports that after a \mathcal{KL}_B analysis of 176 MYCIN rules, he could generate a new knowledge base where 80% of the rules had only a single condition. Further, the problem solving strategies removed all uncontrolled backtracking [9]. However, this result has not been reported elsewhere.

3.2 Are Problem Solving Methods Reusable?

It is not clear that the problem solving methods found by \mathcal{KL}_B are truly reusable. Between the various camps of \mathcal{KL}_B researchers, there is little agreement on the details of the problem solving methods. Contrast the list of inference sub-routines from KADS [46] and SPARK/ BURN/ FIREFIGHTER [21] (termed “knowledge sources” and “mechanism” respectively). While there is some overlap, the lists are different. Also, the number and nature of the problem solving methods is not fixed. Often when a domain is analysed using \mathcal{KL}_B , a new method is induced [20]. Further, different interpretations exist of the same method. For example, the problem solving method proposed by Bredeweg [4] for prediction via qualitative reasoning is different to the qualitative prediction method proposed by Tansley & Hayball [43].

3.3 Do We Need a Model?

This section presents results suggesting that having no model of a problem solving method may be better than having one.

3.3.1 The Corbridge Study

Corbridge *et. al.* report a study in which subjects had to extract knowledge from an expert dialogue using a variety of abstract pattern tools [11]. In that study, subjects were supplied with transcripts of a doctor interviewing a patient. From the transcripts, it was possible to extract 20 respiratory disorders and a total of 304 “knowledge fragments” (e.g. identification of routine tests, non-routine tests, relevant parameters, or complaints).

Subjects were also supplied with one of three problem solving strategies representing models of the diagnostic domain. Each model began with the line “To help you with the task of editing the transcript, here is a model describing a way of classifying knowledge”. Model one was an “epistemological model” that divided knowledge into various control levels of the diagnosis process. Model one was the “straw man”; it was such a vague description of how to do analysis that it should have proved useless. Model two was a KADS problem solving strategy for diagnosis. Model three was “no model”; i.e. no guidance was given to subjects as to how to structure their model. The results are shown in Figure 5.

Model	% disorders identified	% knowledge fragments identified
1 (Epistemological)	50	28
2 (KADS)	55	34
3 (no model)	75	41

Figure 5: Analysis via different models

The statistical analysis performed by Corbridge *et. al.* found a significant difference between the performance of groups 3 compared to groups 1 and 2. Further, no significant difference could be found between the group using the poor problem solving method (model 1) and the group that using a very mature problem solving method (model 2).

These are very counter-intuitive results. Using a hastily-built abstraction (model one) was just as useful as using a mature abstraction (model two). And using no problem solving method worked best of all !! Far from challenging this result, the $\mathcal{K}\mathcal{L}_B$ community is now exploring empirical methods for exploring its approach in the Sisyphus-3 project.

3.3.2 Ripple Down Rules

Compton reports experiments where experts could fix faulty rules using an *unless* patch attached at the end of a rule condition. Patches are themselves rules which can be recursively patched. Experts can never re-organise the tree; they can only continue to patch their patches. These *ripple-down-rule* (RDR) trees are a very low-level representation. Rules cannot assert facts that other rules can use. In no way can a RDR tree be called a model in anything like a $\mathcal{K}\mathcal{L}_B$ sense. Yet this low-level model-less approach has produced large working expert systems in routine daily use. For example, the PIERS system at St. Vincent’s Hospital, Sydney, models 20% of human biochemistry sufficiently well to make diagnoses that are 99% accurate [36]. RDR has succeeded in domains where previous attempts, based on much higher-level constructs, never made it out of the prototype stage [33]. Further, while large expert systems are notoriously hard to maintain [13], the no-model approach of RDR have never encountered maintenance problems. System development blends seamlessly with system maintenance since the only activity that the RDR interface permits is patching faulty rules in the context of the last error. For a 2000-rule RDR system, maintenance is very simple (a total of a few minutes each day).

Compton argues that his process of “patching in the context of error” is a more realistic model of human reasoning than assuming that a human analyst will behave in a perfectly rational way to create some initial correct design [10]. This line is perused further in the *situated cognition* literature where it is claimed that it is folly to use context-independent symbolic assertions to model human reasoning [3, 8, 22, 39]. While many symbolic AI researchers dispute this claim (e.g. [27, 44]), it has some following within the $\mathcal{K}\mathcal{L}_B$ community (e.g. Steels [41] and Clancey [8]).

Our view is not an extreme situated cognition position. We do not reject $\mathcal{K}\mathcal{L}_B$ and KADS out of a disdain for symbolic logics. Like Poole [35], we believe that we can use logics to represent commonsense reasoning in general and context-dependent reasoning in particular, just as long as we don’t use deductive logics. Our preferred approach is *abductive* (see below).

3.4 Is Symbol-Level Modeling Better?

We believe that the high-level abstract view provided by \mathcal{KL}_B obscures important distinctions at the symbol-level. To be fair, the KADS community notes that some overlap exists between problem solving methods: Wielinga *et. al.* note that both monitoring and systematic diagnosis share some processing [46]. However, having noted that some low-level similarities exist between the KL distinctions that they propose, they do not take the next step and simplify their distinctions accordingly. We believe that *abduction* is a symbol-level inference procedure that can unify a wide variety of KL tasks. Consider a dependency and-or graph with invariants \mathcal{I} , edges \mathcal{E} and literals stored in vertices \mathcal{V} . Abduction is the generation of maximal consistent (with respect to \mathcal{I}) worlds \mathcal{W} ($\mathcal{W}_x \subseteq \mathcal{E}$) that contain some subset of the desired goals \mathcal{G} using some subset of known inputs \mathcal{IN} ($\mathcal{IN} \subseteq \mathcal{V}, \mathcal{G} \subseteq \mathcal{V}$). In the case where multiple worlds can be generated, a *BEST* operator selects the preferred world(s). By choosing appropriate \mathcal{IN} , \mathcal{G} , and *BEST*, a range of knowledge base tasks can be implemented. For example:

- Diagnosis [37] uses a *BEST* that favors worlds that explain the most things, with the smallest number of diseases (i.e. maximise $\mathcal{W}_x \cap \mathcal{G}$ and minimise $\mathcal{W}_x \cap \mathcal{IN}$).
- Prediction is implemented by calling the abductive inference with $\mathcal{G} \subseteq \mathcal{V} - \mathcal{IN}$; i.e. find all vertices we can reach from the inputs. Note that in the special case where \mathcal{IN} are all root vertices in the graph and $\mathcal{G} = \mathcal{V} - \mathcal{IN}$, then our abductive system will compute all possible consistent worlds that are extractable from the theory. A more efficient case is that $\mathcal{G} \subset \mathcal{V} - \mathcal{I}$; i.e. some *interesting subset* of the vertices have been identified as possible outputs.
- Validation uses a *BEST* that favours the largest number of covered outputs (i.e. maximise $\mathcal{G} \cap \mathcal{W}_x$) where \mathcal{G} and \mathcal{IN} come from a library of known behaviour of the thing being modeled [26].

Elsewhere, we discuss how to use abduction for classification, explanation, planning, monitoring, qualitative reasoning, verification, multiple-expert knowledge acquisition, explanation, single-user decision support systems and multiple-user decision support systems [24, 23]. We also believe that abduction can model certain interesting features of human cognition [25]. Other authors have discussed the use of abduction for natural-language processing [31], design [34], visual pattern recognition [35], analogical reasoning [15], financial reasoning [18] and machine learning [19].

4 Discussion

Our criticisms of \mathcal{KL}_B should not imply a criticism of \mathcal{KL}_A . When we look into the details of PSCM-SOAR, we find that, like our abductive proposal, PSCM-SOAR uses a low-level uniform structure for KBS design. Further, we take care to distinguish Clancey’s \mathcal{KL}_B work from other \mathcal{KL}_B research. Based on a reverse engineering of previously successful designs, Clancey’s work is founded on documented experience. The same can not be said about other \mathcal{KL}_B work. For example, many of the problem solving methods listed in Tansley & Hayball were created especially for that book by the authors from their own undocumented sources [43, page 260]. Further, Clancey is honest enough to publish his doubts of the symbolic paradigm [8].

References

- [1] H. Akkermans, F. van Harmelen, G. Schreiber, and B. Wielinga. A Formalisation of Knowledge-Level Models for Knowledge Acquisition. In J. Balder and H. Akkermans, editors, *Formal Methods for Knowledge Modeling in the CommonKADS Methodology: A Compilation (KADS-EE/T1.2/TR/ECN/014/1.0)*, pages 53–90. Netherlands Energy Research Foundation, 1992.
- [2] J. Bachant and J. McDermott. R1 Revisited: Four Years in the Trenches. *AI Magazine*, pages 21–32, Fall 1984.
- [3] L. Birnbaum. Rigor Mortis: A Response to Nilsson’s ‘Logic and Artificial Intelligence’. *Artificial Intelligence*, 47:57–77, 1991.
- [4] B. Bredeweg. *Expertise in Qualitative Prediction of Behaviour*. PhD thesis, University of Amsterdam, 1992.
- [5] B.G. Buchanan and R.G. Smith. Fundamentals of Expert Systems. In P.R. Cohen A. Barr and E.A. Feigenbaum, editors, *The Handbook of Artificial Intelligence, Volume 4*, volume 4, pages 149–192. Addison-Wesley, 1989.
- [6] A.N. Campbell, V.F. Hollister, R.O. Duda, and P.E. Hart. Recognition of a Hidden Material Deposit by and Artificially Intelligent Program. *Science*, 217:927–929, 3 September 1982.
- [7] W. Clancey. Heuristic Classification. *Artificial Intelligence*, 27:289–350, 1985.
- [8] W. Clancey. Situated Action: A Neuropsychological Interpretation Response to Vera and Simon. *Cognitive Science*, 17:87–116, 1993.
- [9] W.J. Clancey. Model Construction Operators. *Artificial Intelligence*, 53:1–115, 1992.
- [10] P.J. Compton and R. Jansen. A Philosophical Basis for Knowledge Acquisition. *Knowledge Acquisition*, 2:241–257, 1990.
- [11] C. Corbridge, N.P. Major, and N.R. Shadbolt. Models Exposed: An Empirical Study. In *Proceedings of the 9th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge Based Systems*, 1995.
- [12] R. Davis, H. Shrobe, and P. Szolovits. What is a Knowledge Representation? *AI Magazine*, pages 17–33, Spring 1993.
- [13] A. Van de Brug, J. Bachant, and J. McDermott. The Taming of R1. *IEEE Expert*, pages 33–39, Fall 1986.
- [14] R.O. Duda, P.E. Hart, and R. Reboh. Letter to the editor. *Artificial Intelligence*, 26:359–360, 1985.
- [15] B Falkenhainer. Abduction as Similarity-Driven Explanation. In P. O’Rourke, editor, *Working Notes of the 1990 Spring Symposium on Automated Abduction*, pages 135–139, 1990.
- [16] E. Feigenbaum and P. McCorduck. *The Fifth Generation*. Addison-Wesley, 1983.
- [17] B. Gaines. AAAI 1992 Spring Symposium Series Reports: Cognitive Aspects of Knowledge Acquisition. *AI Magazine*, page 24, Fall 1992.
- [18] W. Hamscher. Explaining Unexpected Financial Results. In P. O’Rourke, editor, *AAAI Spring Symposium on Automated Abduction*, pages 96–100, 1990.
- [19] K. Hirata. A Classification of Abduction: Abduction for Logic Programming. In *Proceedings of the Fourteenth International Machine Learning Workshop, ML-14*, page 16, 1994.
- [20] M. Linster and M. Musen. Use of KADS to Create a Conceptual Model of the ONCOCIN task. *Knowledge Acquisition*, 4:55–88, 1 1992.
- [21] D. Marques, G. Dallemagne, G. Kliner, J. McDermott, and D. Tung. Easy Programming: Empowering People to Build Their own Applications. *IEEE Expert*, pages 16–29, June 1992.
- [22] D. McDermott. A Critique of Pure Reason. *Computational Intelligence*, 3:151–160, 1987.
- [23] T.J. Menzies. An Overview of Abduction as a General Framework for Knowledge-Based Systems. Technical Report TF95-5, Department of Software Development, Monash University, Caulfield East, Melbourne, Australia, 3145, 1995. Available from <http://www.sd.monash.edu.au/~timm/pub/docs/paperonly.html>.
- [24] T.J. Menzies. *Principles for Generalised Testing of Knowledge Bases*. PhD thesis, University of New South Wales, 1995.
- [25] T.J. Menzies. Situated Semantics is a Side-Effect of the Computational Complexity of Abduction. In *Australian Cognitive Science Society, 3rd Conference*, 1995. Available from <http://www.sd.monash.edu.au/~timm/pub/docs/paperonly.html>.

- [26] T.J. Menzies and W. Gambetta. Exhaustive Abduction: A Practical Model Validation Tool. In *ECAI '94 Workshop on Validation of Knowledge-Based Systems*, 1994. Available from <http://www.sd.monash.edu.au/~timm/pub/docs/paperonly.html>.
- [27] M. Minsky. Society of Mind: A Response to Four Reviews. *Artificial Intelligence*, 48:371–396, 1991.
- [28] A. Newell. The Knowledge Level. *Artificial Intelligence*, 18:87–127, 1982.
- [29] A. Newell. Reflections on the Knowledge Level. *Artificial Intelligence*, 59:31–38, February 1993.
- [30] A. Newell, G.R. Yost, J.E Laird, P.S. Rosenbloom, and E. Altmann. Formulating the Problem Space Computational Model. In P.S. Rosenbloom, J.E. Laird, and A. Newell, editors, *The Soar Papers*, volume 2, pages 1321–1359. MIT Press, 1991.
- [31] H.T. Ng and R.J. Mooney. The Role of Coherence in Constructing and Evaluating Abductive Explanations. In *Working Notes of the 1990 Spring Symposium on Automated Abduction*, volume TR 90-32, pages 13–17, 1990.
- [32] K. O'Hara and N. Shadbolt. AI Models as a Variety of Psychological Explanation. In *IJCAI '93*, volume 1, pages 188–193, 1993.
- [33] R.S. Patil, P. Szolovitis, and W.B Schwartz. Causal Understanding of Patient Illness in Medical Diagnosis. In *IJCAI '81*, pages 893–899, 1981.
- [34] D. Poole. Hypo-Deductive Reasoning for Abduction, Default Reasoning, and Design. In P. O'Rourke, editor, *Working Notes of the 1990 Spring Symposium on Automated Abduction.*, volume TR 90-32, pages 106–110, 1990.
- [35] D. Poole. A methodology for using a default and abductive reasoning system. *International Journal of Intelligent Systems*, 5:521–548, 1990.
- [36] P. Preston, G. Edwards, and P. Compton. A 1600 Rule Expert System Without Knowledge Engineers. In J. Leibowitz, editor, *Second World Congress on Expert Systems*, 1993.
- [37] J. Reggia, D.S. Nau, and P.Y Wang. Diagnostic Expert Systems Based on a Set Covering Model. *Int. J. of Man-Machine Studies*, 19(5):437–460, 1983.
- [38] P.S. Rosenbloom, J.E. Laird, and A. Newell. *The SOAR Papers*. The MIT Press, 1993.
- [39] J.R. Searle. Minds, Brain, and Programs. *The Behavioral and Brain Sciences*, 3:417–457, 1980.
- [40] L. Steels. Components of Expertise. *AI Magazine*, 11:29–49, 2 1990.
- [41] L. Steels. How Can we Make Further Progress in Knowledge Acquisition? In R. Mizoguchi, H. Motoda, J. Boose, B. Gaines, and P. Compton, editors, *Proceedings of the Third Japanese Knowledge Acquisition for Knowledge-Based Systems Workshop, JKAW '94*, pages 65–71, 1994.
- [42] M. Stefik, J. Aikins, R. Balzer, J. Benoit, L. Birnbaum, F. Hayes-Roth, and E. Sacerdoti. The Organisation of Expert Systems, A Tutorial. *Artificial Intelligence*, 18:135–127, 1982.
- [43] D.S.W. Tansley and C.C. Hayball. *Knowledge-Based Systems Analysis and Design*. Prentice-Hall, 1993.
- [44] A.H. Vera and H.A. Simon. Situated Action: A Symbolic Interpretation. *Cognitive Science*, 17:7–48, 1993.
- [45] S.M. Weiss, C.A. Kulikowski, and S. Amarel. A Model-Based Method for Computer-Aided Medical Decision-Making. *Artificial Intelligence*, 11, 145-172 1978.
- [46] B.J. Wielinga, A.T. Schreiber, and J.A. Breuker. KADS: a Modeling Approach to Knowledge Engineering. *Knowledge Acquisition*, 4:1–162, 1 1992.
- [47] G. Yost. Implementing the Sisyphus-93 task using SOAR/TAQL. In B.R. Gaines and M. Musen, editors, *Proceedings of the 8th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, pages 46.1–46.22, 1994.
- [48] G.R. Yost and A. Newell. A Problem Space Approach to Expert System Specification. In *IJCAI '89*, pages 621–627, 1989.
- [49] V.L. Yu, L.M. Fagan, S.M. Wraith, W.J. Clancey, A.C. Scott, J.F. Hanigan, R.L. Blum, B.G. Buchanan, and S.N. Cohen. Antimicrobial Selection by a Computer: a Blinded Evaluation by Infectious Disease Experts. *Journal of American Medical Association*, 242:1279–1282, 1979.