

ISSUES WITH META-KNOWLEDGE

TIM MENZIES*, KLAUS-DIETER ALTHOFF†, YANNIS KALFOGLOU‡, ENRICO MOTTA§

*NASA/WVU Software Research Lab, 100 University Drive, Fairmont, WV 26554, USA

†Fraunhofer Institute for Experimental Software Engineering, Kaiserslautern, Germany

‡Institute for Representation and Reasoning; University of Edinburgh, UK

§Knowledge Media Institute, The Open University, Walton Hall, Milton Keynes, UK

E-mail: {tim@menzies.com,klaus-dieter.althoff@iese.fhg.de,yannisk@dai.ed.ac.uk,e.motta@open.ac.uk}

At the SEKE'99 conference, knowledge engineering researchers held a panel on the merits of meta-knowledge (i.e. problem solving methods and ontologies) for the development of knowledge-based systems. The original panel was framed as a debate on the merits of meta-knowledge for knowledge maintenance [21]. However, the debate quickly expanded. In the end, we were really discussing the merits of different technologies for the specification of reusable components for KBS. In this brief article we record some of the lively debate from that panel and the email exchanges it generated.

Keywords: Problem solving methods, ontologies, reuse, knowledge engineering, evaluation.

What is meta-knowledge? How is it being used? Is that usage a useful thing to do? This article approaches these questions from our different viewpoints^a:

The Constructors: Motta and Kalfoglou focus on the *construction* of knowledge base systems (KBS) using two special kinds of meta-knowledge: ontologies [14,33] and problem solving methods (PSMs) [26,6]. PSMs model the useful inference patterns seen in previous applications. Such patterns, it is argued, simplify and clarify future implementations. Ontologies model common domain terminology. This terminology might include the data structures required by a PSM. Using a good ontology, it is argued, can guide developers in the construction of new systems. For more information on ontologies, see <http://www.dai.ed.ac.uk/daidb/people/homes/yannisk/seke99panelhtml.html>.

The Maintainers: Althoff and Menzies focus on the *maintenance* and modification of KBS using case-based reasoning (CBR) [3,1] and continual testing [10,22]. CBR researchers argue that people rarely solve problems from scratch. Rather, they find similar old solutions, then adapt them to new situations. Continual testing is a reflective approach to design; i.e. improvements are prompted by some device reporting errors [30,7]. The most important component of reflective

^aFor more details on these views, see our position papers in the proceedings of SEKE'99.

design is the evaluation device since failed evaluations are essential for improving the design.

Construction vs maintenance is only an approximation of our different positions. Clearly, the two approaches overlap. Constructors believe that maintenance can be simplified via use of good ontologies and problem solving methods. Maintainers believe that their “maintenance” tools can really be used as construction tools since any maintenance task involves some (re)construction of parts of a system.

The following list of reported benefits for meta-knowledge shows specific areas where we agree and disagree. We all agree on the *communications*, *interoperability* and *browsing/searching* benefits. However, while the constructors endorse the *guidance* and *systems engineering* benefits, the maintainers do not.

The communications benefit: Any communication task is simplified by a shared lexicon. Ontologies are a very useful collection of knowledge for systematically viewing and sharing a specific topic/problem. For example, ontologies provide a unified framework within an organization that reduces the terminological confusion [34] arising from different contexts and viewpoints for a particular domain. Meta-knowledge is also useful for any teaching situation. For example, meta-knowledge can simplify a student’s task when (e.g.) reading textbooks. Also, Menzies reports that abstract meta-knowledge descriptions are useful when tutoring software or knowledge engineering [24]. Such abstractions serve as a useful final initiation ritual for a novice designer. When they “get” abstractions, we know that the students are capable of comparing and contrasting a wide range of systems.

The interoperability benefit: Inter-operability among systems with different modeling methods, paradigms, languages and software tools can be achieved with ontologies that act as an inter-lingua [16].

The browsing/searching benefit: The meta-knowledge within an ontology can assist an intelligent search engine with processing a query. For example, if a query returns no results, then the ontology could be used to automatically generalize the query to find nearest partial matches.

The systems engineering benefit: Ontologies and PSMs, it is argued, simplifies system construction. For example, Kalfoglou executes the constraints found in existing ontologies to check new systems. Such pre-existing constraints, are a powerful tool for checking knowledge when other oracles are absent [14]. In other work, one commercial company used the ontology associated with Motta’s PSM design tools to formalize the regulations applicable to the design of a truck cabin. This formalization, associated with a constraint analyzer, cut the design of the geometric layout of the cabin from 4 months to 1 day (!!). Elsewhere, an intelligent PSM librarian was used to build nine KBS applications. Development times changed from one to 17 days (using the librarian) to a range of 63 to 250 days (without using the librarian) [19]. Finally, as a last example, we mention

the SALT KBS editor used for the VT elevator configuration system. SALT restricted its knowledge editors to only those terms relevant for the propose-and-revise PSM used in VT [18,17]. $2130/3062 \approx 70\%$ of VT's rules could be auto-generated by SALT.

The guidance benefit: While we may use little of an ontology or a PSM, it may still be useful as a “pointer tool”. That is, the ontology/PSM could be used as a structuring tool for exploring a new domain. Roughly speaking, reusing abstracted forms of old knowledge is pointing the way saying “these kinds of things are important, even if these particular things are not”. In this approach, developers kick-start the development with an ontology/PSM.

Motta argues strongly for the *systems engineering* and *guidance* benefits. PSM research, he says, has transformed knowledge engineering from “an art” [13] to a structured discipline, organized in terms of a number of generic problem areas (e.g., diagnosis, planning, etc...). The PSM-literate KBS developers can worry less about issues such as (e.g.) conflict resolution strategies for production systems. Such low-level details do not reflect the knowledge-level goals of a domain expert [27,8] In a knowledge-level analysis, the task is to map the space of knowledge engineering techniques and develop handbooks similar to those of other engineering disciplines. These handbooks specify the type of problems we deal with and the kind of techniques we use. In Motta's view PSMs and ontologies provide the key technologies for writing these handbooks.

Menzies disagrees with Motta, arguing that, despite arguments like those made by Motta and Kalfoglou and their colleagues, PSMs are not used widely around the world. Many other approaches exist [20] and some can claim significant successes. For example:

- The United States DARPA High Performance Knowledge Based initiative (HPKB) studied how to write KBs faster [23]. In HPKB year one, the George Mason team generated the most new axioms added per day (787 binary predicates) using DISCIPLE: an incremental knowledge acquisition tool [32]. DISCIPLE includes machine learning tools for abstracting learnt rules which makes them more generally applicable. As DISCIPLE runs, it builds and updates the metaknowledge used for the purposes of abstraction.
- Proponents of randomized search algorithms do not use PSM-style structuring tools for organizing their knowledge bases. Instead, they model their entire domains in simple disjunctive normal form (DNF). In comparative studies, such randomized search over 3DNF has out-performed knowledge-intensive approaches by 2 to 3 orders of magnitude [31].
- In the RDR approach, a KBS stores a patch tree to a knowledge base (each patch fixes a rule and may itself be patched recursively). Candidates for new patches are inferred by an analysis of the paths taken through the KB and the patches found on those paths. Patch histories are low-level syntactic

knowledge yet they capture the context of change of an expert system. Very large expert systems have been built and maintained in this manner, without needing knowledge engineers [10,11,28,29].

Althoff and Menzies have other doubts about systems engineering and the guidance benefits using PSMs and ontologies:

- It may not be cost-effective to routinely devote a significant portion of the systems engineering effort to formalizing most of the domain knowledge. Althoff argues that enshrining domain terminology (e.g. into an ontology) should be the exception, rather than the rule. Over time, the evaluation of some experience base will show which artifacts will be the most useful ones (and only these useful artifacts should be enshrined).
Also, the level of abstraction at which we formalize domain terms should be learnt via extensive experience with that particular term [4]. For example, sometimes, simple text-based descriptions of meta-knowledge are powerful tools. Object-oriented “guidance patterns” serve to direct the analyst’s focus onto a set of issues that previous analysts have found insightful. Such patterns include CHECKS [12], Caterpillar’s Fate [15], and Coad’s strategies [9]. This type of meta-knowledge is stored as simple checklists of English text.
- The systems engineering benefits described above assumed some reuse of ontologies and PSMs. The software engineering experience is that reuse comes at a cost. Verbatim reuse is very rare: often artifacts from a reuse library have to be modified prior to their use. The COCOMO-II software cost estimation model offers an estimate of the cost of adapting reusable sub-routines for a new project [2]. A learning curve must be traversed before a module can be adapted. According to COCOMO-II, by the time you know enough to change a little of that module, you may as well have rewritten 60% of it from scratch (see Fig. 11 in Menzies’ panel notes [21]). The COCOMO-II results relate to the adaption cost of procedural systems. Hence, they may not apply to declarative descriptions of system terminology (i.e. an ontology). However, at the very least, these results caution us that just because we are reusing an ontology or a PSM, this does not necessarily mean that we are building systems more cheaply. Ontologies and PSMs must be learnt prior to use and this learning time may have a non-trivial impact on the overall cost.

It can be argued that that the last point must be false. If it were true, then why can we find the spectacular optimizations of systems development using PSMs and ontologies seen above (recall the examples in the *systems engineering* sections)? Menzies objects to this counter-argument, arguing that without more precise metrics collection, it is hard to interpret these reported optimizations [20,23] (e.g. when Motta’s colleagues reduced their design time from 4 months to 1 day, how much of that reduction was due to the PSM framework and how much to the constraint analyzer that they used?).

Another point of contention was the stability of the meta-knowledge:

- Motta defines meta-knowledge as that knowledge which is stable across domains. Stable knowledge does exist, he says and argues by example: “You do not change the C compiler every time you change the C code”. For many PSMs such as diagnosis: there exist a finite number of different approaches to diagnosis, each one with pros and cons [5].
- Menzies took the opposite view. Unlike Motta, he doubts that any detailed description of meta-knowledge will ever be stable. Meta-knowledge is not useful if it simplifies system construction but complicates system maintenance. Meta-knowledge is still knowledge, he says, and the maintenance of system knowledge is a significant cost of any successful system. Further, he rejects the diagnosis example stating that his reading of the details of the diagnosis literature is that separate and incompatible abstractions are offered by different authors [24]. More generally, based on a recent literature review [25] he claims that in the usual case, expert use of a KBS produces significant changes to that KBS.

In the end, what did we learn? Due to these discussions, we now have a clear definition of where we agree and disagree. Also, we now have an explicit list storing the potential benefits of ontologies and PSMs (to the best of our knowledge, this list was only implicit in the current literature). Lastly, the participants learnt much about each other’s positions which would be useful for future debates (at SEKE 2000?). Based on a flurry of email, we now have the guidelines on what arguments would win over the other side:

- We need more KE metrics. In particular, we need to evaluate the success (or otherwise) of different parts of our systems. Also, we need to test for the stability of our meta-knowledge.
- The constructors/maintainers have to be more explicit about their maintenance/construction methods respectively.

References

1. A. Aamodt and E. Plaza, “Case-based reasoning; Foundational issues, methodological variations, and system approaches,” *AI Communications* **7**(1) (1994) 39–59. Available from http://www.iiia.csic.es/People/enric/AICom_ToC.html.
2. C. Abts, B. Clark, S. Devnani-Chulani, E. Horowitz, R. Madachy, D. Reifer, R. Selby, and B. Steece, COCOMO II Model Definition Manual, Technical report, Center for Software Engineering, USC., 1998. <http://sunset.usc.edu/COCOMOII/cocomox.html#downloads>.
3. K-D. Althoff, A. Birk, S. Hartkopf, W. Muller, M. Nick, D. Surmann, and C. Tautz, “Managing software engineering experience for comprehensive reuse”, in *Proc. 11th Int. Conf. on Software Engineering and Knowledge Engineering, SEKE’99*, Kaiserslauten, Germany, June 1999, pp. 10–19.

4. K.-D. Althoff, M. Nick, and C. Tautz, "Improving organizational memories through user feedback", in *Proc. Workshop on Learning Software Organizations (LSO) (in conjunction with the 11th International Conference on Software Engineering and Knowledge Engineering, SEKE'99, Kaiserslauten, Germany)*, ed. F. Bomarius, June 1999, pp. 27–44.
5. R. Benjamins, *Problem Solving Methods for Diagnosis*, PhD thesis, University of Amsterdam, 1993.
6. R. Benjamins and D. Fensel, "Special issue on problem solving methods", *International Journal of Human Computer Studies* **49**(4) (1998).
7. G. Casady, "Rationale in practice: templates for capturing and applying design expertise", in *Design Rationale: Concepts, Techniques, and Use*, eds. T.P. Moran and J.M. Carroll (Lawrence Erlbaum Associates, 1996), pp. 351–372.
8. W. Clancey, "Heuristic classification", *Artificial Intelligence* **27** (1985) 289–350.
9. P. Coad, D. North, and M. Mayfield, *Object Models: Strategies, Patterns, and Applications* (Prentice Hall, 1997).
10. P. Compton, G. Edwards, A. Srinivasan, P. Malor, P. Preston, B. Kang, and L. Lazarus, "Ripple-down-rules: Turning knowledge acquisition into knowledge maintenance", *Artificial Intelligence in Medicine* **4** (1992) 47–59.
11. P. J. Compton and R. Jansen, "A philosophical basis for knowledge acquisition", *Knowledge Acquisition*, **2** (1990) 241–257.
12. W. Cunningham, "The CHECKS pattern language of information integrity", in *Pattern Languages of Program Design*, eds. J. Coplien and D. Schmidt (Addison-Wesley, 1995). Also available at <http://c2.com/ppr/checks.html>.
13. E. A. Feigenbaum, "The art of artificial intelligence: Themes and case studies of knowledge engineering", in *IJCAI '77*, 1977.
14. Y. Kalfoglou and D. Robertson, "A case study in applying ontologies to augment and reason about the correctness of specifications", in *Proc. 11th International Conference on Software Engineering and Knowledge Engineering, SEKE'99, Kaiserslauten, Germany, June 1999*, pp. 64–71.
15. N. Kerth, "Caterpillar's fate: A pattern language for transformation from analysis to design", in *Pattern Languages of Program Design*, eds. J. Coplien and D. Schmidt (Addison-Wesley, 1995). Also available from <http://c2.com/ppr/catsfate.html>.
16. J. Lee, M. Gruninger, Y. Jin, T. Malone, A. Tate, G. Yost, and other members of the PIF working group, "The PIF process interchange format and framework", *Knowledge Engineering Review* **13**(1) (1998) 91–120.
17. S. Marcus and J. McDermott, "SALT: A knowledge acquisition language for propose-and-revise systems", *Artificial Intelligence* **39** (1989) 1–37.
18. S. Marcus, J. Stout, and J. McDermott, "VT: An expert elevator designer that uses knowledge-based backtracking", *AI Magazine*, Winter 1987, pp. 41–58.
19. D. Marques, G. Dallemagne, G. Kliner, J. McDermott, and D. Tung, "Easy programming: Empowering people to build their own applications", *IEEE Expert*, June 1992, pp. 16–29.
20. T. Menzies, "hQkb — The high quality knowledge base initiative (Sisyphus V: Learning design assessment knowledge)", in *KAW'99: the 12th Workshop on Knowledge Acquisition, Modeling and Management*, 1999. Available from <http://www.tim.menzies.com/99hqkb.pdf>.
21. T. Menzies, "Knowledge maintenance heresies: Meta-knowledge complicates KM", in *11th Annual International Conference on Software Engineering and Knowledge Engineering*, 1999. Available from <http://tim.menzies.com/pdf/99sekekm.pdf>.

22. T. Menzies and C. C. Michael, "Fewer slices of PIE: Optimising mutation testing via abduction", in *SEKE '99*. Available from <http://tim.menzies.com/pdf/99seke.pdf>, 1999.
23. T. Menzies and F. van Harmelen, "Editorial: Evaluating knowledge engineering techniques", *International Journal of Human Computer Studies, Special Issue on Evaluation of Knowledge Engineering Techniques*, 2000. forthcoming; available from <http://tim.menzies.com/pdf/99ekend.pdf>.
24. T. J. Menzies, "OO patterns: Lessons from expert systems", *Software Practice & Experience* **27**(12) (1997) 1457–1478. Available from <http://tim.menzies.com/pdf/97patern.pdf>.
25. T. J. Menzies, "Towards situated knowledge acquisition", *International Journal of Human-Computer Studies* **49** (1998) 867–893. Available from <http://tim.menzies.com/pdf/98ijhcs.pdf>.
26. E. Motta, D. Fensel, M. Gaspari, and A. Benjamins, "Specifications of knowledge components for reuse", in *Proc. SEKE '99*, 1999.
27. A. Newell, "The knowledge level", *Artificial Intelligence* **18** (1982) 87–127.
28. P. Preston, G. Edwards, and P. Compton, "A 1600 rule expert system without knowledge engineers", in *Second World Congress on Expert Systems*, ed. J. Leibowitz, 1993.
29. D. Richards and P. Compton, "Combining formal concept analysis and ripple down rules to support the reuse of knowledge", in *SEKE '97: Proc. 1997 Conf. on Software Eng. & Knowledge Eng.*
30. D. A. Schon, *The Reflective Practitioner* (Harper Collins/Basic Books, 1983).
31. B. Selman, H. Levesque, and D. Mitchell, "A new method for solving hard satisfiability problems", in *AAAI '92* (1992), pp. 440–446.
32. G. Tecuci, "Building intelligent agents: An apprenticeship multistrategy learning theory", *Methodology, Tool and Case Studies* (Academic Press, 1998).
33. M. Uschold and M. Gruninger, "Ontologies: Principles, Methods, and Applications", *The Knowledge Engineering Review* **11**(2) (1996) 93–136.
34. M. Uschold, M. King, S. Moralee, and Y. Zorgios, "The enterprise ontology", *The Knowledge Engineering Review* **13**(1) (1998).