

Validation of Knowledge-Based Systems: Current trends and issues

Alun D. Preece

Department of Computing Science, University of Aberdeen, Scotland, U.K.

Introduction

Assuring the reliability of knowledge-based systems has become an important issue in the development of the knowledge engineering discipline. There has been a workshop devoted to these topics at most of the major AI conferences (IJCAI, AAAI, and ECAI) for the last five years, and the 1994 European Conference on Artificial Intelligence (ECAI-94) in Amsterdam was no exception. The focus of the meeting was on validation techniques for KBS, where *validation* is defined as the process of determining if a KBS meets its users' requirements; implicitly, validation includes *verification*, which is the process of determining if a KBS has been constructed to comply with certain formally-specified properties, such as consistency and irredundancy.

The Amsterdam workshop was an intimate meeting, and the fifteen attendees were predominantly from European institutions. In spite of - or perhaps because of - this intimacy, the workshop succeeded in highlighting many of the significant trends and issues within its area of concern. The purpose of this short article is to review the trends and issues in question, drawing upon the contributions made during the workshop.

Topics and Trends in KBS Validation

The formal contributions to the workshop programme were clustered around three general themes: theoretical foundations of KBS validation, KBS validation in practice, and experimental techniques for KBS validation. These contributions were supplemented by an invited session on the use of formal specification languages in KBS validation, and a closing discussion centring around challenges and significant unsolved problems in the area.

In recent years, the main technological theme in the KBS validation area has been the development of tools for automatic verification of knowledge bases. Within this sub-area, the dominant concern has been with the "first generation" type of rule-based systems, and the verification has been aimed at detecting *anomalies* - such as subsumed or conflicting rules - which are symptomatic of logical faults in the knowledge base (Preece et al., 1992). This aspect of KBS validation technology has now become reasonably mature, reflected by the fact that none of the presentations at the Amsterdam workshop were directly concerned with it.

Several aspects of current work does build upon the earlier approaches, however. For example, a proposal to extend the principles of anomaly detection to nonmonotonic knowledge-based systems was made by Zlatareva (1994), and is one of the first attempts to address this type of system. A second example is the work of Wendler and Vignollet (1994), who are working towards applying a coherence-checking technique, originally developed for conventional monolithic KBS, to modular knowledge bases. These works would seem to be motivated by the changing nature of the KBS field itself, as it moves away from the "first generation" rule-based expert systems, towards knowledge-based systems with more sophisticated architectures and reasoning mechanisms.

If we regard the previous trend essentially as "technological pull" from KBS technology, then we can also see a "technological push" acting upon the validation area from AI technology in general. Examples of this can be seen in two papers at the Amsterdam workshop. Menzies and Gambetta (1994) successfully exploited a technique developed from abductive reasoning, in which an exhaustive abduction procedure is used to determine if a qualitative model is consistent with given test cases. Bouali et al. (1994) adapted Reiter's diagnosis theory to perform correction of faulty KBS, considering different means of doing so depending upon whether reliable test cases are available. This last work is an instance of a class of work which is gaining prominence in the validation area: that of integrated approaches to automatic validation *and fault rectification*.

Issues Arising in KBS Validation

It is probably reasonable to assert that the most significant issue currently facing European researchers in the KBS validation area is of how to exploit the work being done in the area of formal specification languages (FSL). In recent years several languages have been developed which allow knowledge engineers to create an initial, formal description of a KBS, and progressively refine it towards an implementation (Fensel and van Harmelen (1994) provide a good introduction to this area). Clearly, FSL can play a significant role in KBS validation.

In an invited talk, Frank van Harmelen argued how one FSL, (ML)², can be used to bridge the gap between an informal description of a KBS, and a completely formal description, in such a way as to provide a degree of assurance that the formal description reliably captures the users' requirements. In this way, validation is "built-into" a rigorous process of specification and refinement, so that the final product will have been validated without requiring testing. A subsequent paper by Weusten (1994) on the development of KBS in the legal profession reinforced this view of validation; here, KBS were generated semi-automatically starting from formal specifications in the form of decision trees and tables; because the domain experts can easily understand the initial specifications, they have confidence in the final product, and therefore do not require it to be tested in any other way. Checks for consistency, irredundancy and other properties are incorporated into the KBS-generation procedure.

Complementary viewpoints on the use of FSL for KBS validation were provided by other presenters. Treur and Willems (1994) are developing criteria against which they would wish to verify KBS specified in their FSL DESIRE. Vermesan and Wergeland (1994) are using algebraic specification techniques derived from conventional software engineering, to develop a KBS in the domain of market analysis in shipping. An open question arising from these works is that of whether FSL designed specifically for KBS-like software are more appropriate than "wide spectrum" FSL from software engineering, for supporting KBS validation.

Where do we Go from Here?

In this short article, we have touched upon some of the themes which are becoming increasingly important in the area of KBS validation. There are important unanswered questions surrounding these themes, questions which will likely motivate much of the immediate future work in the area. To conclude, here is a (far from complete) summary of themes and unsolved problems.

Formal specification for KBS validation: As yet, there is no satisfactory proposal as to how formal specification techniques can be used effectively with validation techniques, and vice versa. We lack methods by which formal specifications can be utilised fully in performing validation and verification. We also lack methods by which the specifications can themselves be validated and verified.

Validation and refinement of KBS: While there has been some success in employing techniques from other areas of AI (including machine learning, theory refinement, diagnosis) to provide a way not only to find flaws in a KBS, but to rectify them also, this technology is rather limited at present. In particular, we know a lot more about refining static domain knowledge than we do about refining dynamic control knowledge.

Validation of new kinds of KBS: While the task of verifying rule-based systems for logical properties is now well-understood, comparatively little work has been done validation methods for other kinds of KBS. Areas of concern which have been identified include: distributed KBS (for example, co-operative distributed problem-solvers), hybrid KBS (built using a combination of different programming paradigms), non-monotonic KBS, and model-based KBS. It would seem that, while some validation issues are common across different types of system, some issues are quite different. A good example is the notion of consistency-of-knowledge: this is usually desirable in a monotonic rule-based system, but would we require consistency between the knowledge of different agents in a distributed KBS? This is arguable.

Given these open-ended considerations, it would appear that the ECAI-94 workshop on validation of KBS will not be the last such meeting!

Acknowledgements

As chairman of the workshop described in this article, I would like to thank all of the participants for making it a success, the workshop programme committee for their efforts, and the organisers of ECAI-94 for being such excellent hosts. This article reflects my personal view of the results of the workshop.

References

Bouali, F, Loiseau, S and Rousset , MC, 1994. "KBS Correction: a proposal based on diagnosis theory". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 21-29.

Fensel, D and van Harmelen, F, 1994. "A comparison of languages which operationalise and formalise KADS models of expertise". *Knowledge Engineering Review*, **9**.

Menzies, T and Gambetta, W, 1994. "Exhaustive abduction: a practical model validation tool". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 70-76.

Preece, A, Shinghal, R and Batarekh A, 1992. "Principles and practice in verifying rule-based systems". *Knowledge Engineering Review*, **7**, 115-141.

Treur, J and Willems, M, 1994. "On verification in compositional knowledge-based systems". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 4-20.

Vermesan, A and Wergeland, T, 1994. "A formal approach to knowledge base refinement and verification". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 42-52.

Wendler, B and Vignollet, L, 1994. "Verification of the coherence of modular knowledge bases". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 77-85.

Weusten, M, 1994. "Validation and verification in legal knowledge based systems". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 53-69.

Zlatareva, N, 1994. "Checking non-monotonic knowledge bases for inconsistencies and incompleteness". *Proceedings of the ECAI-94 Workshop on Validation of Knowledge-based Systems*, pp 30-40.