AI Magazine Volume 19 Number 3 (1998) (© AAAI)

Verification and Validation of Knowledge-Based Systems

Report on Two 1997 Events

Grigoris Antoniou, Frank van Harmelen, Robert Plant, and Jan Vanthienen

This article gives an overview of two recent events on the validation and verification of knowledge-based systems: (1) the 1997 European Symposium on the Verification and Validation of Knowledge-Based Systems (EUROVAV'97) and (2) the Fourteenth National Conference on Artificial Intelligence Workshop on the Verification and Validation of Knowledge-Based Systems. To give an integrated view of current research issues in this field, we organized this article along thematic lines, unifying the reports of the two separate meetings. Our report focuses on the trends that we think will be important in the near future in this field.

This article gives an overview of two recent events on the validation and verification of knowledge-based systems: (1) the 1997 European Symposium on the Verification and Validation of Knowledge-Based Systems (EUROVAV'97) and (2) the Fourteenth National Conference on Artificial Intelligence Workshop on the Verification and Validation of Knowledge-Based Systems. To give an integrated view of current research issues in this field, we organized this article along thematic lines, unifying the reports of the two separate meetings. Our report focuses on the trends that we think will be important in the near future in this field.

The 1997 edition of EUROVAV was already the fourth time that the symposium was held. It was chaired by Jan Vanthienen (University of Leuven, Belgium) and Frank van Harmelen (Vrije Universiteit, Amsterdam) and held in the beautiful city of Leuven, Belgium. With 25 submissions (of which 16 were accepted) and 35 attendants, EUROVAV'97 was roughly the same size as other recent meetings. Detailed information on the program, the participants, and online abstracts of all the papers can be found at www.econ.kuleuven.ac.be/congres/eurovav/eurovav97.htm.

Annual meetings on the verification and validation of knowledge-based systems have been organized in North America since 1988. The 1997 event took place in Providence, Rhode Island. It was chaired by Grigoris Antoniou (Griffith University, Australia) and Robert Plant (University of Miami) and attracted about 20 participants from academia and industry.

Discussion of Selected Topics

In this section, we discuss the following topics: rule-based systems, formal methods, abstract knowledge models, modular knowledge bases, machine-learning techniques, multiagent systems, uncertainty, and web site tools.

Rule-Based Systems

Rule-based systems are still the most widely used representation method for the development of knowledge-based systems. In addition, they have been around for many years, so it comes as no surprise that they have been studied extensively in the validation and verification community. Early work focused on redundancy, consistency, circularity, and so on.

One basic approach in the study of rule-based systems is the representation of the knowledge using matrixes, which are then analyzed and manipulated to detect several anomalies. However, this approach has suffered from drawbacks with respect to both the storage space and the processing power required. In their EUROVAV'97 paper entitled "Verification and Validation of Rule Bases Using a Binary Encoded Incidence Matrix Technique," Coenen and Dunne address these problems by introducing a binary representation of incident matrixes. Their approach leads to reduced storage use and supports more efficient processing of the matrixes using logical operations on the bit level.

In their AAAI-97 validation and verification paper "Performance Assessment and Incremental Evaluation of Rule-Based Systems," Chander, Shinghal, and Radhakrishnan argue that as a knowledge-based system evolves, its evaluation should be incremental because redoing validation and verification work from scratch can be costly. Their technique is based on providing a link that connects the conceptual, design, and implementation levels of a system. Knowledge is acquired using goal specification to capture the problem-solving states. Based on the specified goals, the structure of the system at the implementation level is defined by a set of rule sequences that infer goals.

Formal Methods

In software engineering, the use of formal specification languages for the purposes of verification has long been advocated. In knowledge engineering, formal methods have not played a prominent role until now. Recent years have seen a gradual increase in the use of formal methods in knowledge engineering, and they are also finding their way toward validation and verification. Fensel and Schonegge's EUROVAV'97 paper "Specifying and Verifying Knowledge-Based Systems with KIV" reports on their work on formally specifying a knowledge-based system using the Karlsruhe interactive verifier (KIV), which was originally constructed for specifying classical programs. With KIV, the structure and contents of a knowledge base are specified using predicate logic, but the inference strategy of the knowledge-based systems is expressed in dynamic logic. This approach allows them to prove properties such as termination and correctness of the knowledge-based system with the assistance of the KIV theorem prover.

Quite similar in aim, but rather different in approach, was the EUROVAV'97 paper by Cornelissen, Jonker, and Treur, "Compositional Verification of Knowledge-Based Systems: A Case Study in Diagnostic Reasoning." This paper proposes using structural properties of a knowledge-based system as the main guidance during verification. In particular, the authors concentrate on the hierarchical decomposition of the system into subcomponents as the main structure for their proofs. Their approach is to split the proof of a required property into a number of lemmas, each of which can be proved separately. The contribution of the authors lies in the guidelines they give for organizing the required set of lemmas, that is, by (1) strictly following the hierarchical decomposition of the system, (2) only formulating lemmas in terms of subcomponents of one level deeper in the decomposition hierarchy, and (3) formulating lemmas that only concern a single component.

Abstract Knowledge Models

In the past decade, a gap has appeared between two areas within knowledge engineering: (1) knowledge acquisition and (2) knowledge validation. The knowledge-acquisition community has developed methods for modeling expertise in forms that are still close to the original knowledge as formulated by the expert and, therefore, that can still be some way removed from an efficiently implementable form. The verification community, however, has mostly concentrated on efficient implementation forms such as production rules. Two EUROVAV'97 papers aimed at closing this gap by trying to exploit the high-level knowledge-acquisition models for the purposes of verification. Both papers used KADS expertise models as the basis for verification, albeit for different application areas. In their paper entitled "Validation and Verification of Diagnostic Systems," Van Harmelen and Ten Teije showed how a general model of diagnostic reasoning could be used to prove properties of a knowledge-based system based on such a model. Marcos, Moisan, and Pobil in "A Model-Based Approach to the Verification of Program Supervision Systems" showed similar results for a knowledge-based system that performed program supervision (that is, the automatic configuration and execution of existing pieces of software for a given task). Of course, such high-level models can only be used as the basis of verification if they are sufficiently precise in both form and content; so, it is not surprising that both these papers were closely connected with the use of formal specification methods for the knowledge-based system, as mentioned previously.

Modular Knowledge Bases

When knowledge bases are large, it is not possible to verify all the knowledge simultaneously. In such cases, it is necessary to decompose the knowledge bases into smaller partitions to perform verification. Two EUROVAV '97 papers proposed different approaches to this problem: Ramaswamy and Sarkar ("Global Verification of Knowledge-Based Systems via Local Verification of Partitions") used directed hypergraphs as a formalism to structure a knowledge-based system into partitions. Essentially, each hypernode corresponds to a (set of) clause used in conditions or conclusions of rules, and hyperlinks correspond to sets of rules that make up the inference path from one hypernode to another. This graph is then partitioned on the basis of a precedence ordering among the nodes (closely related to the dependency graph between the rules). Verification can then proceed by locally verifying individual partitions in the graph and subsequently propagating these local verification results through the dependency graph.

Whereas Ramaswamy and Sarkar used rules as their knowledge representation formalism, the paper by Vanthienen, Mues, and Wets, entitled "Intertabular Verification in an Interactive Environment," represents knowledge in tabular form; separate tables form the modularization mechanism in their case. However, one of the main drawbacks of these systems is that anomalies that occur because of the interaction of tables are neglected. Their paper investigates an approach to dealing with these so-called intertabular anomalies. An interesting aspect of their work is the strong requirement that their techniques must be used online by a knowledge engineer during knowledge-based system development. As a result, their approach uses heuristics in those cases where exhaustive checks would be too inefficient.

A second EUROVAV'97 paper dealing with modularity of tabular rule bases was "Feature Construction for Verification and Validation of Tabular Knowledge Base Systems" by Piramuthu. His aim was to find techniques that would help in breaking up large tables into smaller and modular tables that are more easy to deal with. Machine-learning techniques for automatic feature construction were used for this purpose.

Machine-Learning Techniques

Verification of a knowledge-based system must often be followed by an action to repair or improve the knowledge-based system. It has long been recognized that there is a close relation between such knowledge base repair actions and machine learning. In his excellent invited lecture at EUROVAV'97, Luc De Raedt showed how techniques from inductive logic programming (currently one of the most active and successful machine-learning techniques) could be related to the validation and verification of knowledge-based systems.

A number of contributed papers at EUROVAV'97 also emphasized the connection with machine learning: In their paper "Inductive Hypothesis Validation and Bias Selection in Unsupervised Learning," Talavera and Cortes proposed a framework for automatic validation of machine-induced knowledge-based systems based on the capability of shifting the bias in the inductive learning system. Their experimental results could be seen as a contribution to both validation and verification and machine learning.

The paper entitled "The Selection of Training Cases for Automated Knowledge Refinement" by Palmer and Craw considered the role of training cases in knowledge-based system refinement. They showed how the choice of training cases has a crucial effect on the quality of the refined knowledge base and argued that it is therefore unacceptable to select training cases at random. They described how training cases can specially be selected to validate specific refined knowledge bases, and they show how this selection leads to a higher quality of the knowledge base but uses a smaller number of training cases.

Multiagent Systems

Multiagent systems are a recent development in AI. They arose out of the realization that to solve certain kinds of problem, it is useful to develop a system in which a number of knowledge-based systems cooperate and combine their problem-solving capabilities. Each knowledge-based system is constructed as a software agent that has autonomy and interacts with other agents to solve the given problem.

In his AAAI-97 validation and verification paper "Verification of Multiagent Knowledge-Based Systems," O'Leary gives an overview of correctness and verification issues for multiagent systems. His work concentrated mainly on issues of interagent verification. In particular, it studies cases where verification is conducted on a metarule base generated from the rules in each of the agents' knowledge bases. In these cases, existing knowledge base verification tools can be applied. Anomalies such as conflicts, circularity, subsumption, inconsistency, and completeness are studied in this context. Also, the paper identifies the problem of agent isolation.

Although O'Leary's paper takes a fairly abstract view, in their AAAI-97 validation and verification paper "Verifying Multiagent Knowledge-Based Systems Using COVERAGE," Preece and Lamb study a specific agent approach, the ARCHON architecture. This architecture distinguishes between two layers: (1) the ARCHON layer, which is concerned with interaction and communication and (2) an intelligent system layer that contains the knowledge of a particular agent and can be a knowledge-based system. Within this architecture, Preece and Lamb study multiagent systems that are realized as rule-based systems. They consider three types of knowledge: (1) domain knowledge, which corresponds to the intelligent system layer; (2) cooperation knowledge, which corresponds to the ARCHON layer; and (3) a monitoring unit, which defines the links between the two layers. The paper studies anomalies that can occur in the interaction of these different kinds of knowledge. It describes the COVERAGE tool that detects these anomalies and is an extension of the COVER tool that validates classical rule-based systems. The paper concentrates mainly on intraagent anomalies (that is, cooperation knowledge–domain knowledge anomalies), but work on cooperation knowledge–domain knowledge anomalies is in progress.

Uncertainty

From the development of the earliest knowledge-based systems, it has been recognized that it can be necessary to represent and reason with uncertainty of the knowledge used. Approaches to uncertainty reasoning include probability theory, certainty factors, the Dempster-Shafer theory, and fuzzy logic. At the AAAI-97 validation and verification workshop, there were two papers dealing with correctness and verification issues related to Bayesian networks:

In their paper "BVAL: Probabilistic Knowledge-Base Validation," Santos, Gleason, and Banks discuss the validation of Bayesian knowledge bases (BKBs), a generalization of classical Bayesian networks capable of incorporating more detailed probabilistic dependencies and incomplete knowledge of states and dependencies. Their approach requires human involvement in the validation process. The authors argue that in the face of incomplete information, a major feature of BKBs, it is unrealistic to expect that validation will be conducted fully automatically. Rather, a human needs to interact with the system during validation to correct errors stemming from incomplete information by modifying the knowledge in the knowledge base. The aim of their tool BVAL is to minimize the interaction required.

In "MACK: A Tool for Acquiring Consistent Knowledge under Uncertainty," Santos, Banks, and Banks describe a tool for knowledge acquisition and maintenance in the framework of BKBs. The tool, MACK, guarantees the consistency of the information stored in a BKB, both as it is acquired and later as it is maintained. The system performs incremental checks and reports to the expert's inconsistencies (both logical and probabilistic validity ones). MACK has been applied to the National Aeronautics and Space Administration posttest diagnostic system that supports main engine analysis for the Space Shuttle.

Web Site Tools

Although the World Wide Web is constantly gaining importance, the tools supporting the web still lack satisfactory functioning. For example, information retrieval is still based on syntactic criteria using search engines, but semantic search would be desirable. Another common problem are links pointing to nonexistent pages.

It is the latter problem that Rousset analyzes in her EUROVAV-97 paper "Verifying the World Wide Web: A Position Statement." She argues that this problem is mainly caused by updates, and it can be addressed by adding more structure into web pages. Although this solution is impractical for the entire World Wide Web, Rousset argues that it can be done on the level of web sites, collections of web pages under a common administration. The paper discusses how modeling the semantics of a web site can be used to define integrity constraints that prevent the introduction of anomalies when the content of web pages is updated.

Conclusion

In our opinion, the field of validation and verification of knowledge-based systems is in a state of transition. The problems that have traditionally been studied are now well understood, with feasible algorithms and a theoretical underpinning. Although the uptake of these techniques by industry still leaves a lot to be desired, the academic community is now moving on to new problems. In particular, the use of formal specification techniques, the use of more abstract knowledge models, and new application areas such as multiagent systems and the World Wide Web were prominent new directions at both of the events described here.

We expect these themes to play an important role in the 1998 meetings on validation and verification of knowledge-based systems, including a European meeting held in conjunction with the Sixth International Conference on Principles of Knowledge Representation and Reasoning and an American validation and verification meeting held as a AAAI-98 workshop. The next EUROVAV symposium will be held in 1999 in Oslo, organized by Alun Preece and Anca Vermesan.

Grigoris Antoniou is an associate professor in Computing and Information Technology at Griffith University, Australia. His research interests include verification and validation of knowledge systems and knowledge representation and reasoning and their application in software engineering and information systems. His e-mail address is ga@cit.gu.edu.au.

Jan Vanthienen is professor of information systems at Katholieke Universiteit Leuven (Belgium), Department of Applied Economic Sciences, Information Systems Group. His current research interests include knowledge-based systems and decision tables, knowledge representation, verification and validation, information systems analysis and design, help desk automation, and computer-based training. He has published more than 70 refereed papers in international journals and conference proceedings. In addition to serving as a program committee member on various conferences and workshops, he was the program cochair of the 1997 European Symposium on the Validation and Verification of Knowledge-Based Systems. His e-mail address is frankh@cs.vu.nl.

Robert T. Plant is an associate professor in the Department of Computer Information Systems at the University of Miami, Coral Gables, Florida. Plant obtained his Ph.D. in computer science at The University of Liverpool, England. Previously having studied at The Programming Research Group, Oxford University, England, and Wadham College, Oxford. Plant was the chairman of the workshop on Validation and Verification of Knowledge-Based Systems at the Twelfth National Conference on Artificial Intelligence and cochair of the 1997 workshop. Plant is a chartered engineer (U.K.), a European engineer, and a fellow of the British Computer Society and holds a visiting professorship in computer science at the University of Wolverhampton in England and was a visiting associate at Templeton College, Oxford, England, in 1997. His research interests are in virtual organization, knowledge-based systems, and software engineering. His e-mail address is rplant@exchange.sba.miami.edu

Frank van Harmelen is a lecturer at the Free University of Amsterdam, where he is working on formal specification languages for knowledge-based systems and their use in validation and verification. From 1990 to 1995, he was working on the KADS Project at the University of Amsterdam, where he contributed to the development of the ML^2 language for formally specifying knowledge-based systems. He is author of a book on metalevel inference and editor of a book on knowledge-based systems and has published over 30 research papers. His e-mail address is frankh@cs.vu.nl.