The Knowledge Engineering Review, Vol. 15:2, 2000, 187–196
Printed in the United Kingdom

Validation and verification of knowledge-based systems: report on EUROVAV99

FRANS COENEN[1], TREVOR BENCH-CAPON[1], ROBIN BOSWELL[2], JULIETTE DIBIE-BARTHÉLEMY[3], BARRY EAGLESTONE[4], RIK GERRITS[5], ERIC GRÉGOIRE[6], ANTONI LIGĘZA[7], LUIS LAITA[8], MIECZYSLAW OWOC[9], FLORENCE SELLINI[10], SILVIE SPREEUWENBERG[5], JAN VANTHIENEN[11], ANCA VERMESAN[12] and NIRMALIE WIRATUNGA[2]

[1]*Department of Computer Science, University of Liverpool, UK*
[2]*School of Computer & Mathematical Science, The Robert Gordon University, Aberdeen, UK*
[3]*Department of OMIP, INA-PG, Universite Paris IX—Dauphine, France*
[4]*Department of Information Studies, University of Sheffield, UK*
[5]*LibRt B.V., Postbus 90359, 1006 BJ Amsterdam, Netherlands*
[6]*CRIL, University of Artois, France*
[7]*Institute of Automatics, AGH, Krakow, Poland*
[8]*Department of AI, Universidad Politecnica de Madrid, Spain*
[9]*Wroclaw University of Economics, Wroclaw, Poland*
[10]*PSA Peugeot Citroen, La Garenne Colombes, France*
[11]*Department of Applied Economic Science, Katholieke Universiteit Leuven, Belgium*
[12]*Det Norska Veritas, Oslo, Norway*

1 Introduction

Knowledge-Based (KB) technology is being applied to complex problem solving and safety and business critical tasks in many application domains. Concerns have naturally arisen as to the dependability of Knowledge-Based Systems (KBS). As with any software, attention to quality and safety must be paid throughout development of a KBS, and rigorous Verification and Validation (V&V) techniques must be employed. Research in V&V of KBSs has emerged as a distinct field only in the last decade, and is intended to address issues associated with quality and safety aspects of KBSs, and to provide such applications with the same degree of dependability as conventional applications. In recent years, V&V of KBSs has been the topic of annual workshops associated with the main AI conferences, such as AAAI, IJCAI and ECAI.

In this paper we report on EUROVAV99, the *Fifth European Symposium on Verification and Validation of Knowledge Based Systems and Components (EUROVAV99*, `http://www.dnv.no/research/safekbs/eurovav99/`), which was held in Oslo in the summer of 1999, and was sponsored by Det Norske Veritas and the British Computer Society's Specialist Group on Expert Systems (SGES). EUROVAV is the leading European conference dedicated to the advancement of the theory and practice of V&V of KBSs. Its objective is to bring together researchers from both academia and industry, not only from Europe, but also from all over the world, encouraging all styles of V&V approaches and a variety of application areas. EUROVAV99 was the fifth in a biannual series, with previous meetings in Cambridge, UK (1991), Palma de Mallorca, Spain (1993), Chambery, France, (1995), and Leuven, Belgium (1997).

The symposium successfully integrated two contrasting streams. On the one side, the theoretical stream presented the best of recent developments in V&V of KBSs, and their components, including database systems, covering a wide range of technical areas. On the other side, the applications stream presented papers describing practical applications of KBSs, and the associated V&V approaches. The symposium was opened by Professor Pal Bergan, director of research for Det Norske Veritas (DNV).

For each stream there were keynote addresses by acknowledged international experts in their field. For the theoretical stream, Trevor Bench-Capon, from University of Liverpool, addressed the

role of ontologies in the Verification and Validation of KBSs. His presentation provided an excellent starting point for EUROVAV99, since the use of ontologies proved to be a theme that re-occurred throughout the symposium. For the application stream there were two keynote speakers: the first was Jeffrey Voas from the USA, who talked about user participation-based software certification; the second was an invited speaker from DNV, Per Martinsen, who talked about the certification activities of software/knowledge-based components integrated within programmable systems. The presentations were all followed by lively and fruitful discussions.

The symposium dealt with all aspects of KB V&V, but broadly, the presentations could be classified as follows:

- Theory and techniques
- Ontologies
- Safety critical KBS (the Safe-KBS project)
- Knowledge Revision and Refinement
- Applications
- Certification
- The wider picture.

A brief synopsis of each of the papers presented under each of these headings is presented in the following sections.

2 Theory and techniques

In this first section, five contributions were presented. The first was intended to carry out a certain amount of *scene setting*, while the remaining four were concerned with more detailed issues. The scene setting paper was given by Mieczyslaw Owoc of the Wroclaw University of Economics (Owoc et al., 1999), and provided an overview/insight into one aspect of KBS validation, i.e., the knowledge element. The authors presented a slightly unorthodox view of KBS and V&V, with one considered to be a sub-component of the other. Owoc's principal concern was that of validating knowledge (rather than the structure of a KBS). In this regard, the authors suggest a number of guiding principles, dividing them into the more global set apart from specific guidelines oriented on the particular form of a knowledge base.

Eric Grégoire of the Universite d'Artois was concerned with the validation of KBSs, concentrating on the possibility of the KB containing contradictory knowledge (Grégoire et al., 1999). Grégoire noted that the checking of the logical consistency of rulebases expressed in a clausal propositional logic form is a NP complete problem (see Cook (1971) or Garey and Johnson (1979) for proofs). To tackle this problem, Grégoire suggested that consistency checking should be performed incrementally during the development of a KB rather than on completion, and then proposed a mechanism to achieve this based on a depth-limited instantiation schema. Grégoire developed this technique within a nonmonotonic representation formalism, and exhibited various forms of inconsistencies that are linked to the progressive nature of his instantiation schema. The technique is implemented by means of a powerful and non-standard use of local search techniques, allowing very large rulebases to be checked for consistency and inconsistency.

The contribution by Luis Laita from the Universidad Politecnica de Madrid (and his research team) dealt with the automated extraction of knowledge and verification of consistency in rule-based KBSs (Laita et al., 1999). To achieve this objective, they have defined an algebraic model that can be applied to many-valued logics, which is incorporated into a software package written in the computer algebra language CoCoA. Luis Laita commenced by summarising the main idea of modelling propositional many valued logic with modal operators (interpreted over Lukasiewicz matrices) using polynomial residue class rings over an appropriate field (of integers modulo p, where p is the number of logical values). The crucial consequence of this construction is that tautological consequences and consistency can be expressed by the properties of ideals of the corresponding polynomial residue class rings. Given a KBS the system translates the rules and facts contained in

polynomials. When dealing with these polynomials, to be a tautological consequence is equivalent to an ideal membership problem (that can be decided by checking whether a "Normal Form" is 0 or not), and the inconsistency of the KBS is equivalent to the degeneracy of an ideal (that can be decided by checking whether a certain "Groebner Basis" is 1 or not).

Silvie Spreeuwenberg and Rik Gerrits from LibRT Redeployment Technology (Netherlands) presented a tool (Spreeuwenberg and Gerrits, 1999) to verify a KB using meta-rules, to be used within the object-orientated KBS development environment Aion8 from Platinum Technology. Besides OO, Aion8 supports decision tables, rule-based programming, introspection, code generation and Component Based Development (CBD). (For further details on Aion, see `http://www.platinum.com`, and search for *Aion 8* or *AionDS*.) The V&V tool is designed to be used both during and after the development of KBSs, and focuses on the logical verification of such systems. The tool uses meta-rules, meta information and the inference engine of Aion to accomplish this task. By using the same inference engine in the verification process as the inference engine used to evaluate and fire the rules to be verified, Spreeuwenberg and Gerrits suggest that there can be no discrepancy between the run time logic and the logic used in the validation process. This first version is developed in and for Aion, but the method and algorithms are applicable to any environment that supports rule-based programming with an inference engine.

Much of the work described above is concerned with the standard rule-based representation of knowledge. Juliette Dibie-Barthélemy of the University of Paris IX Dauphine, however, considered the semantic validation of conceptual graphs (Sowa, 1984) used to represent knowledge based on constraints given by experts for validation purposes (Dibie-Barthélemy et al., 1999). Using this approach a KB is semantically valid if it satisfies all the given constraints (the authors use the term *specification* to describe expert knowledge given for validation purposes).

3 Ontologies

The ontologies concept, as applied to KBSs, is a fairly recent invention. Essentially, an ontology in the computer science sense of the word, is a set of definitions of (hierarchically ordered) classes, objects, attributes, relations, and/or constraints (axioms) whose main function is to provide a well defined vocabulary for the expression of domain knowledge. Thus, an ontology is a knowledge-level description (Newell, 1982), in that it is independent of any representational formalism (Van Heijst, 1985). Also, an ontology is considered to be a meta-level description, because it is a specification of a specification. It is often said that an ontology describes the domain knowledge that remains invariant over various knowledge bases in a certain domain (cf. Guarino and Giaretta, 1995).

In general, we can say that ontologies may contribute to the following five KBS development areas:

1. Domain-theory development.
2. Knowledge acquisition.
3. Knowledge-system design.
4. System documentation.
5. Knowledge exchange.

The role that the ontologies concept has to play in the evaluation of KBSs was examined by the first of the invited speakers, Trevor Bench-Capon of the University of Liverpool (Bench-Capon and Jones, 1999). The first half of the presentation was concerned with the motivation of ontologies—"an explicit specification of the conceptualisation of a domain". Bench-Capon emphasised that an ontology was not a model, "an ontology defines classes of concepts while models include instances of concepts" (Gruber, 1993). Bench-Capon then went on to state that ontologies describing a KBS can be validated; and introduced a software tool, PRONTO, to support this approach to KBS evaluation. PRONTO operates by using the ontology to generate test cases which are then executed against the current knowledge base. If the test case fails, depending on the nature of the failure, and input from the expert, either the ontology or the knowledge base is redefined.

The second talk in this category was given by Robert Engels of CognIT (Norway) whose presentation was "CORPORUM—An Automated Ontology Construction Tool". Robert Engels' talk focussed on the need to support knowledge companies who are in the business of capturing knowledge for Knowledge Management (KM). The talk included a demonstration of the CORPORUM tool, which used the notion of an ontology to extract knowledge.

The final presentation in this category was by Henrik Giaever, of DNV, was entitled "Verification and Validation in Support of Knowledge Management". This was a very theoretical presentation concentrating on a philosophical discussion on the nature of knowledge, especially the differing views of information (consciousness). Perhaps the most interesting feature of the talk was Giaever's augmentation of the well known distinction between *knowing how* or *knowing that* by *knowing why*.

4 Safety-critical KBS (the safe KBS project)

The two presentations in this section described the work carried out as part of the Safe-KBS ESPRIT project, the aim of which is to define development and certification methodologies for KBSs, especially embedded KBSs used in safety-critical domains. The particular result of this work has been the definition of a Safe-KBS lifecycle, designed so that it can be integrated with formal methods and provide appropriate support for software certification methods.

The first presentation in this section (Dondossola, 1999) was by Giovanna Dondossola from ENEL-SRI (Milan), and started by providing an overview of the results of the "Safe-KBS" project and the role of formal methods within that project. This was followed by a detailed discussion on the requirements for KBS formal methods within the context of both specification and V&V. Dondossola then continued to consider the impact of such methods on the KBS lifecycle, pointing out that the lifecycle models should be defined independently of specific development methods and techniques, and then going on to consider the elements of such a lifecycle based on that developed as part of the Safe-KBS project. Finally, Dondossola gave consideration to the selection and evaluation of KBS methodologies, concluding that no optimum formal method exists appropriate to the full spectrum of KBS development requirements.

The contribution by Rune Steinberg of Computas (Norway) also described results founded on the Safe-KBS project, but concentrating on KBS design issues (Steinberg et al., 1999). Steinberg argued that, despite the fact that a well designed lifecycle will contribute substantially to the confidence in a KBS, there is also a particular need to concentrate on design. Steinberg then went on to show how the Safe-KBS lifecycle is supported by a catalogue of design patterns for safety-critical KBSs which describe design solutions that seek to increase the confidence of using KBS technology. The nature of these design patterns was therefore the main focus of this contribution.

5 Knowledge revision and refinement

Having established that errors/anomalies exist in a KBS system, the practitioner will wish to revise/maintain the KBS in question. Given that it is possible automatically to detect errors in rule-bases, the obvious next step is the automated refinement of such rule-bases. Tools to achieve this end are usually labelled as *knowledge refinement* or *knowledge revision* tools. Such tools are typically used to assist knowledge engineers by automatically identifying the locations within a KB where the knowledge may be need to be changed, and (importantly) suggesting appropriate amendments (Craw, 1998).

The first contribution in this section, by Robin Boswell of the Robert Gordon University Aberdeen, described a suite of KB refinement tools under development as part of the KRUST-Works project (Boswell and Craw, 1999; Craw and Boswell, 1999). These have been applied successfully to both artificial and real-world KBSs written in a variety of different languages. Central to this toolkit is a set of generic refinement operators and a hierarchical representation language for KBS rules. This knowledge hierarchy constitutes an ontology for KBS *rule-elements*; each KBS rule condition and conclusion is made up of one or more such elements. The ontology

captures knowledge about the behaviour of rule-elements that is invariant over different shells and environments. Its purpose is to organise the refinement operators and apply them appropriately, in the same way that Fensel et al. (1997) use ontologies to organise problem solving methods and apply them to appropriate tasks.

The contribution by Nirmalie Wiratunga, also of the Robert Gordon University, looked at search techniques which enable KRUSTWorks to perform iterative knowledge refinement of faulty rule-bases. Incorporating backtracking into a best-first search solves the problem of getting caught in local maxima (Wiratunga and Craw, 1999). Next, a re-ordering strategy can be used to make the best use of examples which when first encountered are correctly solved by the faulty rule-base, but may provide fault evidence for a subsequent modified rule-base. Finally, the search effort can be reduced by identifying pairs of examples providing fault evidence that result in conflicting repairs, and re-ordering them next to each other so that such examples are processed in successive iterations.

The second contribution in this section, by Stephan Loiseau of the Universite d'Angers, was a discussion of the distinction/overlap between KBS validation and KBS revision/refinement—i.e., maintenance (Dupin de Saint-Cyr and Loiseau, 1999). This is a continuation of similar work studying the overlap between KBS verification and revision (see Bouali et al., 1997a, 1997b). In the presentation, Loiseau defined refinement as a modification of an incoherent KB, to restore its coherency. Revision in turn is the modification of an inconsistent KB to restore consistency. A coherent rule-base is one which, if and only if with any valid input, contradictory results cannot be inferred.

Jan Vanthienen, of the Katholieke University of Leuven, described a modularisation approach to restructuring KBSs founded on a decision table technique (Vanthienen and Moreno-Garcia, 1999). Vanthienen is an internationally recognised expert on decision table techniques, especially in the context of KBSs. Modularisation is an important concept with respect to verification, validation and maintenance, as it encapsulates knowledge into isolated modules which are easier to verify and maintain. By encapsulating knowledge and by avoiding redundant specification of knowledge, the risk of introducing incompleteness or inconsistencies during maintenance is minimised. The paper is illustrated with a specific KBS application (SEAR—Enterprise Human Resources Advisory System), the aim of which is to optimise human resource management in a social and juridical consultantship setting based on the Spanish Social Security General Regime Law. Vanthienen used his approach and the Prologa decision table tool to restructure the SEAR KBS into a decision table hierarchy, while dealing with intratabular and intertabular verification.

6 Applications

Much of the foregoing concentrates on the theory of KBS V&V. An important aspect of EUROVAV is also the practical application of V&V techniques.

To this end, Francisco Ribeiro, of the Instituto Superior Tecnico in Lisbon, presented a discussion of V&V of a KBS system for supporting human experts in assessing application for the house renovation grant system in Portugal (Ribeiro, 1999). Ribeiro began by considering the development of the system (called HRGS), and then went on to consider the particular approach taken to verify and validate this system.

The contribution in this section by Jorge Santos, of the Polytechnic Institute of Porto, focused on the V&V of intelligent applications software used by the electricity supply industry (Santos et al., 1999). In particular, the presentation addressed the V&V of a KBS (called SPARSE) used to assist operators of Portuguese Transmission Control Centres in incident analysis and power restoration. Although this was the central theme, the contribution also provided a more general insight into the application of KBS V&V techniques.

Florence Sellini of PSA Peugeot Citroen, presented a discussion of the verification of product modelling for design aid systems for mechanical sets based on the knowledge of expertise (Sellini and Yvars, 1999). The talk described how knowledge used in a design system can be verified. A particular knowledge representation was described, and additional constraints formulated which

should conform to this knowledge. The presentation made a useful contribution as an example of what is/can be/must be done in a realistic KBS application.

The final contribution in this section (Schroeder, 1999) was concerned with the formal verification of workflows, and was by Michael Schroeder of City University London. Schroeder addressed the problem of formal verification of business processes. In the past few years, much research has been devoted to formal verification of business processes (van der Aalst, 1998; Wodtke and Weikum, 1997). As motivation, Wodtke and Weikum point out that the only validation techniques for workflows currently available are animation and simulation, and argue that symbolic verification is a mature field that provides a number of results and tools appropriate for workflow verification.

Guided by an example business process for a BT customer correspondence centre, in which work items are distributed, collected, possibly fixed, and then dispatched, Michael Schroeder showed how to translate a given process specification in the standardised PIF format into the calculus for concurrent systems CCS. The main difficulty in the translation arises as PIF is under-specified, and does not distinguish whether an activity has to wait for all preceding activities to finish or just for one. Schroeder's translation commits to the former interpretation based on practical considerations for the application addressed.

Once the formal specification in CCS is given, one can use verification tools to check for a variety of interesting properties such as deadlocks, livelocks, liveness and safety properties such as, for example, ensuring that jobs are sometimes fixed and always dispatched. Schroeder also addressed shortcomings of the tools, such as the inability to detect livelocks directly.

7 Certification

The certification of software and associated techniques is an area that is rapidly expanding. Much of the knowledge that has been developed over the past few years is now mature enough to be implemented in industry (Vermesam, 1997). In this section, three papers were presented which dealt with software certification.

Jeffrey Voas (invited speaker), in the first paper (Voas, 1999), noted that the most popular approaches to software certification are founded on publisher oaths and auditors (who carry out spot-checks to confirm such oaths), and suggests that this approach may be flawed—*Dirty water can run from clean pipes*. Voas went on to propose a certification methodology that does not employ auditors and publisher oaths, i.e., independent product certification. Voas is a recognised expert in the field of software assessment, having developed, amongst other things, the theory of inoculating programs against error as a way to assess the safety of critical software systems.

The second talk, also invited, was presented by Per Martinsen, principal surveyor and head of the "Control Systems and Nautical Safety" department at DNV, who talked about "Software certification activities for programmable systems". The talk described the certification process in the wider context, especially with respect to the practice of DNV.

The third paper in this section, presented by Anca Vermesan, DNV Norway, presented an approach to certifying software contained within systems where its functionality is considered critical within the environment in which it is placed (Vermesan et al., 1999). Two certification methods were described in detail, and the role of V&V in certification of KB components emphasised. Vermesan argued that, as techniques for the evaluation of software systems develop in parallel with the systems themselves, certification bodies must constantly review the requirements that are required from software components in order to achieve safe and dependable systems. The approach presented here is also based on the findings of the Safe-KBS European project.

8 The wider picture

In this final section, some wider aspects of KBS V&V were considered, in particular:

1. The possible application of KBS V&V to recent database models.

2. The contribution that integrity checking techniques current within the database community may have to play with respect to KBS V&V.
3. The application of conventional software V&V techniques to KBS V&V.

Frans Coenen, of the University of Liverpool, considered possible future directions for rule-based V&V on the presumption that many techniques developed by the rule-based V&V community have come out of the research domain into what might be described as the main stream (Bench-Capon et al., 1998). In particular Coenen considered opportunities for possible cross-overs between database integrity checking techniques and a number of rule-based V&V strands (Coenen et al., 1998). Collectively, they refer to these techniques as VVI (Validation Verification and Integrity) techniques. The presentation began by giving an overview of integrity issues, then went on to consider overlaps between database and rule-based systems, and finished with a discussion of future directions in VVI.

Continuing in this theme, the second paper in this section, presented by Antoni Ligęza of the Institute of Automatics of the AGH in Cracow, Poland, was concerned with the analysis and verification of Data and Knowledge (D&K) using a common model founded on an extension of the relational database model (Ligęza, 1999). The model basically takes the tabular form based on attributes; appropriate logical interpretation allows for common model for both data and knowledge. Similarities and difference between data and knowledge item were discussed, and it was claimed that within the model data can be considered to be part of the knowledge expressed with the finest granularity and unconditional.

The main focus of the paper was the meta-level analysis of D&K tables to identify and define problems associated with adequate data and knowledge representation. In the paper, seven top-level problems were identified; the proposed taxonomy considered the following: adequacy of D&K representation, sufficiency, constraints satisfaction and consistency, correctness, equivalency and manipulation (an "algebra" of KBS), similarity and analogy, and classification (including inductive generalisation). Logical definitions were proposed within a unifying logic-based perspective. Influence of selected issues on system reliability, safety and quality was analysed and a uniform logical framework for discussion was put forward. An important final note was that, although the problems above are specified with logical definitions, for the simplified, tabular D&K model they can be approached with algebraic methods, which seems to move them closer to engineering applications. Ligęza also went on to consider various technical problems encountered during the implementation of verification procedures.

The third paper in this section was by Barry Eaglestone of the University of Sheffield, described some of the V&V issues of concern in the database community in the context of schema integration of federated databases (Masood and Eaglestone, 1999). Eaglestone also briefly overviewed a semi-automated methodology for the schema analysis phase of schema integration based upon the ontology-based conceptualisation of schema semantics. In this approach, semantics of schema elements are made explicit by mapping them into an ontology, as an aid to verifying correctness of schema integrations and the validity of schema designs.

The final paper in this section was by Frode Hgberg, also of DNV Norway, who considered the application of conventional software V&V techniques to KBS V&V (Vermesan and Høgberg, 1999). Høgberg presented an experimental framework, founded on the Safe-KBS project, to determine whether a technique is applicable or not, based on concepts found in mutation testing. The framework itself consisted of a number of steps guiding the researcher/practitioner in the assessment process. Mutation testing is used to simulate faults in an example program to determine the techniques ability to detect them. The paper concludes that the framework gives basis for qualitative assessment of the applicability and efficiency of applying specific traditional V&V and testing techniques to KBS components. To illustrate these ideas Vermesan and Hgberg seeded a sample KBS with faults through a mutation process, and then used a number of traditional V&V methods to find these faults.

9 Conclusions

The main themes of the EUROVAV99 conference were:

- The role of ontologies in KBS V&V and their components.
- Certification of KBS.
- The V&V of KBS founded on the object-orientated paradigm.
- V&V of multiple KBS.
- The effect that KBS V&V "know how" might have on other disciplines.

As with previous meetings, EUROVAV99 provided a valuable opportunity to keep up to date with developments in V&V of KBS, and to share experiences in the practical issues of developing and certifying knowledge-based systems in particular, and software systems in general.

The proceedings are published by Kluwer Academic (Vermesan and Coenen, 1999). Anyone wishing to obtain a copy should contact Anca Vermesan (Anca.Vermensan@dnv.com) at DNV, Oslo. Plans are currently underway for EUROVAV 2001.

References

Bench-Capon, TJM, Castelli, D, Coenen, FP, Devendeville-Brisoux, L, Eaglestone, B, Fiddian, N, Gray, A, Ligęza, A and Vermesan, A, 1999. "Validation, verification and integrity issues in expert and database systems" *Expert Update* **2**(1) 31–35.

Bench-Capon, T and Jones, D, 1999. "PRONTO: Ontology-based evaluation of knowledge-based systems" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Boswell, R and Craw, S, 1999. "Organising knowledge refinement operators" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Bouali, F, Loiseau, S and Rousset, M-C, 1997a. "Verification and revision of rules bases" In: Hunt, J and Miles, R, editors *Research and Development in Expert Systems XIV, Proceedings of ES97*, SGES Publications, pp. 253–276.

Bouali, F, Loiseau, S and Rousset, M-C, 1997b. "Revision of rule-bases" *Proceedings EUROVAV 97* Katholieke Universiteit Leuven, Belgium, pp 193–203.

Coenen, FP and Bench-Capon, TJM, 1993. *Maintenance of Knowledge Based Systems: Theory, Tools and Techniques*. Academic Press, London.

Coenen, F, Eaglestone, B and Ridley, M, 1999. "Validation, verification and integrity in knowledge and data base systems: future directions" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Cook, SA, 1971. "Complexity of theorem-proving procedures" *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing* pp 151–158.

Craw, S, 1998. "Krust works: Developing a generic refinement toolkit" *Expert Update* **1**(2) 35–47.

Craw, S and Boswell, R, 1999. "Representing problem-solving for knowledge refinement" *Proceedings of the Sixteenth National Conference on Artificial Intelligence* in press.

Dibie-Barthélemy, J, Haemmerlé, O and Loiseau, S, 1999. "Constraint for validation of conceptual graphs" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Dondossola, G, 1999. "Formal methods for the engineering and certification of safety-critical knowledge based systems" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Dupin de Saint-Cyr, F and Loiseau, S, 1999. "Validation and refinement versus revision" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Fensel, D, 1995. "Formal specification languages in knowledge and software engineering" *The Knowledge Engineering Review* **9**(4).

Fensel, D, Motta, E, Decker, S and Zdrahal, Z, 1997. "Using ontologies for defining tasks, problem-solving methods and their mappings" In: Plaza, E and Benjamins, R, editors *Knowledge Acquisition, Modeling and Management, Proceedings of the 10th European Workshop (EKAW '97)*, Sant Feliu de Guixols, Spain. Springer-Verlag, pp 113–128.

Garey, MR and Johnson, DS, 1979. *Computers and Intractability – A Guide to the Theory of NP-Completeness*. W. H. Freeman and Co., San Francisco.

Grégoire, E, 1999, "Progressive instantiation for the logical validation of nonmonotonic KBs" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Gruber, T, 1993. "Toward principles for the design of ontologies used for knowledge sharing" *Technical report KSL-93-4*, Knowledge Systems Laboratory, Stanford University, Stanford, USA.

Guarino, N and Giaretta, P, 1995. "Ontologies and knowledge bases: towards a terminological clarification" In: Mars, NJI, editor *Towards Very Large Knowledge Bases*. IOS Press.

van Harmelen, F and ten Teije, A, 1997. "Validation and verification of conceptual models of diagnosis" *Proceedings EUROVAV97*, Katholieke Universiteit Leuven, Belgium.

van Heijst, G, 1995. The role of ontologies in knowledge engineering, *Doctoral Thesis*, University of Amsterdam, Amsterdam, Netherlands.

Laita, L, Rianes-Lozano, E, Maojo, V and de Ledesma, L, 1999. "Computer algebra based verification and knowledge extraction in RBS: Application to medical fitness criteria" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Ligęza, A, 1999. "Intelligent Data and knowledge analysis and verification: Towards a taxonomy of some specific problems" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Masood, N and Eaglestone, B, 1999. "Ontology-based verification and validation of federated database systems" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Newell, A, 1982. "The knowledge level" *Artificial Intelligence* **18** 87–127.

Owoc, M, Ochmańska, M and Gladysz, T, 1999. "On principals of knowledge validation" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Riberio, F, 1999. "Certification and validation of a multistrategy knowledge based system" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Santos, J, Ramos, C, Vale, Z and Marques, A, 1999. "Validation and verification of knowledge-based systems for power system control centres" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Schroeder , M, 1999. "Verification of business processes for a correspondence handling center using CCS" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Sellini, F and Yvars, PA, 1999. "A priori verification of product models in mechanical design" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Sowa, JF, 1984. *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley.

Spreeuwenberg, S and Gerrits, R, 1999. "A knowledge base tool to validate and verify an Aion KB" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Steinberg, R, Fjellheim, R and Olsen, S, 1999. "Design patterns for safety-critical knowledge based systems, knowledge revision and refinement" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

van der Aalst, WMP, 1998. "The application of Petri nets to workflow management" *Journal of Circuits, Systems and Computers*.

Vanthienen, J and Moreno-Garcia, A, 1999. "Illustrating knowledge base restructuring and verification in a real world application" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Vermesan, A, 1997. "Quality assessment of knowledge-based software: Certification considerations" *Proceedings of Third IEEE International Software Engineering Standards Symposium (ISESS '97)* Walnut Creek, CA.

Vermesan, A and Coenen, F, 1999. *Validation and Verification of Knowledge Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Vermesan, A and Høgberg, F, 1999. "Applicability of conventional software verification and validation to knowledge based components: A qualitative assessment" In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Vermesan, A, Sjøvaag, J, Martinsen, P and Bell, K, 1999. "Verification and validation in support for software

certification methods'' In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Voas, J, 1999. ''User participation-based software certification'' In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.

Wodtke, D and Weikum, G, 1997. ''A formal foundation for distributed workflow execution based on state charts'' *Proceedings of the International Conference on Database Theory*.

Wiratunga, N and Craw, S, 1999. ''Incorporating backtracking search with knowledge refinement'' In: Vermesan, A and Coenen, F, editors *Validation and Verification of Knowledge-Based Systems: Theory, Tools and Practice*. Kluwer Academic, Utrecht, Netherlands.