

Formal methods in knowledge engineering

FRANK VAN HARMELEN¹ and DIETER FENSEL²

¹*Vrije Universiteit Amsterdam, Department of Mathematics and Computer Science, AI Group, De Boelelaan 1080a, 1081 HV Amsterdam, The Netherlands (email: frankh@cs.vu.nl)*

²*SWI, University of Amsterdam, Roetersstraat 15, 1018MT Amsterdam, The Netherlands*

Abstract

This paper presents a general discussion of the role of formal methods in knowledge engineering. We give an historical account of the development of the field of knowledge engineering towards the use of formal methods. Subsequently, we discuss the pros and cons of formal methods. We do this by summarising the proclaimed advantages, and by arguing against some of the commonly heard objections against formal methods. We briefly summarise the current state of the art and discuss the most important directions that future research in this field should take. This paper presents a general setting for the other contributions in this issue of the journal, which each deal with a specific issue in more detail.

1 Historical growth of knowledge engineering towards formal methods

Although the history of KBS technology and knowledge engineering (KE) is well documented in a number of places in the literature (e.g., Jackson, 1990, Chapt. 2), in this section we will give an account of the development of KE* which will show the natural growth of this field to the use of formal methods. In the development of KE methods and technology, we distinguish three main periods: the programming period; the modelling period; and the current move towards the introduction of formal methods.

The first era of KE technology stretched from the late 1970s to the mid 1980s. This period was characterised by the development of new programming techniques. New systems were described in terms of the representation techniques that they employed; rules, frames, Horn clauses, semantic networks, etc. KBS development environments gave support at the level of these representation techniques, and often aimed at integrating these different representations (e.g. ART, KEE, Knowledge Craft; see Laurent et al., 1986 for a comparison). Such programming techniques were often developed and widely used before a proper formal understanding of them was available (witness, for example, the certainty factor model in MYCIN (Buchanan & Shortliffe, 1984) or multiple inheritance in LOOPS (Keene, 1989)). Although shells like EMYCIN (van Melle, 1979) and others made some progress towards abstraction, they were still presented in terms of the representation language and inference engine they provided.

The move away from this first period was already heralded as early as 1980 by Alan Newell (1982) in his “knowledge level” lecture, (1982) but should more properly be situated as late as 1985 when Clancey (1985) published his “heuristic diagnosis” paper. In this paper, Clancey analysed a number of systems at a higher level of abstraction than simply their code. In particular, he identified a specific problem solving method which formed the basis of the behaviour of a number of systems, even though they were all coded in different ways. Clancey called this method “heuristic classification”, and described it in terms of its essential inference steps and the types of

*At least from a European perspective. It is well possible that from an American or Japanese perspective, a different picture would emerge.

knowledge manipulated by these inference steps. Most importantly, this analysis was entirely independent from the particular way this method could be programmed in any particular representation language. This type of analysis triggered the development of a number of knowledge engineering methodologies. The late 1980s saw a number of such methodologies which were all aimed at so called “knowledge level” analysis of KBS tasks and domains. Generic Tasks (Chandrasekaran, 1986), KADS (Wielinga & Breuker, 1986; Weilinga et al., 1992), Methods-to-Tasks (Musen, 1989) and Role-limiting methods (Marcus, 1988) are some of the prominent examples of such methods.

These approaches all differ on the structure that they propose for analysing knowledge, on the degree of task-specificity, on their link with executable code, and many other properties, but all of them are based on the idea of constructing a “conceptual model” of a system which describes the required knowledge and strategies at a sufficiently high level of abstraction, independent from any particular implementation formalism.

Although these methods were highly successful and widely adopted (e.g. Ueberreiter & Voß, 1991; Bauer & Karbach, 1992; Löckenhoff et al., 1993 for KADS), almost all of them are often criticised from both within and from outside the KBS community (e.g. from software engineering) for their informality and corresponding lack of precision. Even though many of these methods claimed to be based on Newell’s knowledge level hypothesis, for which Newell himself had proposed logic as the ideal language for analysis, none of the most prominent methods in the later 1980s incorporated much formal analysis. At best, these methods offered structured and semi-formal notation without clear semantics (e.g., the inference structures of KADS), or they interfaced directly with executable inference and representation mechanisms (e.g., Generic Tasks), but none provided either a formal syntax or a mathematical semantics, let alone formal derivation rules.

This lack of formal backbone in most if not all of the leading KE methodologies was apparently felt in many places, since the third period in this historical sketch, beginning in the early 1990s, brought a plethora of attempts at formalising much of the work that had been done before. The arguments that were used to motivate these formalisations will be discussed in more depth in section 3. They concerned issues like the removal of ambiguity, the possibility of formally deriving properties like soundness or completeness, and bridging the gap between informal conceptual model and design of the executable system. Much of the formal language development took place in Europe, and because of the prominence of the KADS conceptual model in European KE, a large number of different proposals for formulating KADS were published: MODEL-K (Karbach et al., 1991), OMOS (Linster, 1993), MoMo (Voss & Voss, 1993; Gebhardt, 1995) (all from GMD in Bonn), FORKADS (Wetter & Schmidt, 1991; Wetter, 1990) from IBM Heidelberg, (ML)² (van Harmelen & Balder, 1992) from Amsterdam. Other languages were not directly based on KADS, but were based on a conceptual model closely related to KADS: KARL (Angele et al., 1995; Fensel, 1995) from Karlsruhe based on the MIKE model (Angele et al., 1993), K_{BS}SF (Veld et al., 1994) from the Dutch Telecom based on the VITAL model (Jonker et al., 1991), and TSL (Pierrot-Goldbreich & Talon, 1995) from Paris based on the TASK model. Yet other languages were not originally based on KADS, but it was shown how they could be used for KADS-like models. QIL (Shadbolt et al., 1992) from Nottingham, GCLA (Aronsson et al., 1994) from SICS in Sweden, AIDE (Kassel & Grebovol, 1993) from Compeigne and MODEL (Barbuceanu, 1993). KADS was certainly not the only source of formal languages. Formal modelling languages based on other conceptual models are DESIRE (van Langevelde et al., 1992) from the Free University of Amsterdam, MC/GETFOL (Giunchiglia et al., 1993) from Trento, and MILORD (Sierra & Godo, 1993) from Blanes. Full references and brief comparisons of these languages will be given in section 3. All of these languages aim at capturing the prominent distinctions made in these conceptual models: distinguishing different knowledge types, describing the use of these knowledge types in inference steps, a specification of control knowledge, and aiming at task- and domain-independent descriptions. Other work applied existing formal specification languages to KBS systems or architectures. Examples include the use of Z (Craig, 1991a,b), CCS (Velthuisen, 1992;

van Melle, 1979) and OBJ3 (Nakagawa et al., 1993). The community has been actively organising itself: workshops at ECAI'92 (Vienna) and ECAI'94 (Amsterdam), as well as regular informal meetings (Bonn '92, Karlsruhe '93, Bonn '94, Amsterdam '95 and Paris '96), plus a WWW resource at <ftp://swi.psy.uva.nl/pub/kem1/kem1.html>. The work in this field has been extensively reviewed (Fensel & van Harmelen, 1994) for KADS-based languages and (Treur & Wetter, 1993) for a collection of languages based on a variety of conceptual models. These reviews will be summarised in section 3.

The evolution of KE as sketched in these three periods also has direct consequences for the life-cycle employed in KBS development. In the first period, knowledge acquisition was seen as a direct transfer of human problem-solving expertise to a computer program. The acquired knowledge was immediately represented by a running prototype. In the second period, knowledge acquisition is viewed not as a transfer process but as a modelling activity. The result of knowledge acquisition is no longer only a running program, but a conceptual model that describes problem-solving expertise in an implementation independent manner. As a consequence, a large gap arose between the outcome of knowledge acquisition and the final implementation of a KBS. Conceptual models described in natural language are insufficient input for the implementation step, because it is mainly a question of natural language interpretation whether an implementation fulfils such a specification. The intuition of a programmer has to fill in the gaps and has to resolve ambiguity in such models. Normally it is not at all clear whether a programmer has the necessary domain and task knowledge to do this properly. The development of formal modelling techniques were a natural answer to this shortcoming by defining an intermediate level between semi-formal models and implementations. They enable a precise specification of the KBS and the required knowledge, while still abstracting in a twofold manner from the implementation. First, they abstract from implementational details (e.g., whether a set will be implemented by a list or by an array). Secondly, an implementation is normally optimised to improve its efficiency. Such efficiency aspects which are not related to the expertise but to the way it is implemented by appropriate data-structures and algorithms are of no concern during knowledge acquisition. Using a formal specification which is not used as an efficient implementation of the KBS enables to abstract from these efficiency aspects which are only related to its implementation. Using the implementation of a KBS as its specification leads either to a mix-up of these quite different aspects or leads to non-efficient implementations which can at best be useful in prototyping. Therefore, formal specification techniques save the abstraction from the implementational level as achieved by the knowledge-level hypothesis of Newel but add substance and formal preciseness to it.

It seems fair to say that the knowledge level hypothesis of Newel was a necessary step as it enabled to get rid of implementational details when discussing models of expertise. Otherwise, it caused significant irritation in its original formulation by Newel. Debates arose about what could or should be specified at the knowledge level, and about appropriate formalisms (if any). This irritation is not at all a surprise as it could not be clear from the beginning what the appropriate notion for specifying KBS's would be. Meanwhile, a number of approaches exist which characterise what should be specified at the knowledge level and what are the appropriate modelling primitives to express such models. As a consequence, work on formal specification techniques could be done which defines formal semantics for these new modelling primitives. In that way the semiformal specifications at the knowledge level can be viewed as a necessary *intermediate* step. The knowledge-level hypothesis enabled the research community to escape from the limitations of the available implementation formalisms and it created the *possibility* and the *necessity* for new types of formal specification techniques. It enforced the development of formal specification techniques which reintegrate the precision and unambiguity of implementation formalisms into knowledge level specifications without losing their conceptual structure and without getting confused by implementational aspects.

The purpose of this paper is to review and discuss the motivations for this body of recent work (section 2), to briefly survey the state of the art (section 3), and to identify future research directions (section 4). Another paper in this issue deals with the relation of the formal methods in

KE with their counterparts in software engineering and information systems. The final paper in this issue discusses the potential of these formal approaches for validation and verification of KBS.

2 Pros and cons of formal methods in knowledge engineering

In software engineering, the usefulness of formal methods has been a hotly debated topic during many decades. During the recent growth of knowledge engineering towards the use of formal methods, many of the same debates resurface in the knowledge engineering community (e.g., Fensel, 1994). In this section, we will first very briefly repeat the acclaimed advantages for formal methods. We will take the advantages that are claimed for formal methods in software engineering (see, e.g., Wing, 1990), and reinterpret these in the context of knowledge engineering. In the second and major part of this section, we will follow Hall (1990) and Bowen and Hinchey (1994) by listing some of the most prominent myths about the use of formal methods in knowledge engineering (KE) in particular, but also in software engineering (SE) in general. We try to dispel most if not all of these myths.

2.1 Advantages of formal methods in knowledge engineering

The advantages of the use of formal methods can be distinguished by the phase of the system development to which they contribute:

Modelling: The conceptual models that are constructed in a modern knowledge engineering process are typically defined using a combination of natural language and graphical elements. As a result, these models are defined only informally or at best semi-formally. As is well known, such documents in natural language have an ambiguous and imprecise semantics. There are as many meanings for a specification as there are readers, and it is a question of text interpretation as to whether a specification is as intended for a system. The ambiguity of such conceptual models in a knowledge engineering context was aptly illustrated by Aben's analysis of the use of a single inference step from the KADS framework in a number of papers in Bauer and Karbach (1992). Aben (1995, p. 36) states that out of nine papers that use the abstract inference step, only one author uses it in correspondence with the original definition in Breuker et al. (1987). Most of the other papers use different versions of the inference step, even including versions with different numbers of input arguments. Models which are expressed in a formal specification language have a precise semantics grounded in mathematical representations and can disambiguate the informal representations. In principle, direct implementations of the informal models can also serve this purpose. However, this would imply either committing the conceptual model to particular machine-oriented details (in order to obtain an efficient implementation) or to ignore such details with an efficient and badly designed implementation as a result (see also myth 1 below).

Design: Although formal specifications can often themselves not be directly efficiently executed, they can form a bridge between the informal high-level conceptual models and detailed machine-oriented design. This is because formal specifications do provide additional detail to the informal models, but still do so at a high conceptual level. An approach of the use of formal representation as a bridge between conceptual models and detailed design for KBS can be found in Landes and Studer (1995).

Furthermore, formal specification languages open the possibility for gradual and stepwise refinement of a specification towards an efficient implementation, where each of the refinement steps is provably correct. This then guarantees the correctness of the final implementation with respect to the initial specification. Advances in this area have been made particularly with the VDM and CIP languages (Bicarregui et al., 1994).

Evaluation and maintenance: Although not much exploited in practice yet (neither in software engineering nor in knowledge engineering) the use of a system description which is at the same time precise and at a high level of abstraction would be a great asset during the maintenance of a system, provided the specification of the system would be kept up to date with the evolving implementation.

Besides these advantages related to the various phases of system development, the following points apply to each of the phases mentioned above:

Validation and verification: Because formal specifications are mathematical objects, they can be subjected to formal manipulation and proofs. This can (at least in principle) lead to proofs of desired properties such as completeness and soundness of a specification with respect to the requirements. Informally, completeness states that the entire required I/O relation can indeed be derived by the specification, while soundness states the converse: the specified I/O does not exceed the required relation. When subjected to formal analysis, these concepts break up into more refined versions (e.g. Treuer & Willems, 1994). Besides these two properties, other properties are candidates for formal verification and validation, such as termination, consistency, irredundancy, etc. For a discussion on this issue, see also the accompanying paper by Preece and Meseguer in this issue of the Journal (1995).

Besides formal proofs, validation and verification may also take place through testing if the formal specification is executable. KBS specification languages like KARL, TSL, DESIRE and others have both a solid formal foundation and are executable (even if not efficiently). Particularly in the absence of proof techniques that scale up to specifications of realistic size, this use of "specifications as prototypes" is the only currently available practical path today to the use of formal specifications in validation and verification.

Reuse of components: Many of the knowledge engineering methodologies mentioned in section 1 place great emphasis on the reuse of various modelling components, varying from the reuse of standardised primitive inference steps (e.g. Aben, 1995, which formalises a set of standard steps proposed in Breuker et al., 1987), to the reuse of complex problem solving methods through libraries of predefined elements (e.g., Breuker & van de Velde, 1994). As discussed above, Aben (1995) showed the problems in reusing informally defined primitive inference steps, and experiences in the Sisyphus II project show that the situation is similar with regard to problem solving methods which are only defined informally. Only names are reused, but they refer to inferences or problem solving methods with totally different or even contradictory meanings. Development effort and costs could be noticeably reduced by the actual reuse of predefined elements with a standardised meaning which comes from their formal definition. This does not necessarily have to mean that the entire reuse process is formalised: Fensel (1994,b) is an example of an informal analysis of the development of a family of systems for design-tasks which has benefited from the availability of formally defined components out of which the systems were constructed (see also the discussion under myth 5 below).

2.2 False objections to formal methods in knowledge engineering

2.2.1 Myth 1: The soft knowledge of experts cannot be adequately formalised

If this claim were true, it would not only be the end of formal methods in KE, but of KE as an engineering discipline overall. After all, if it were true that soft expert-knowledge cannot be formalised, it would also apply, and to a much greater extent, to the final implementation of a KBS, because an executable computer program has for fundamental reasons a more limited expressive power than, e.g., first order predicate logic. In addition, the natural language description of the expertise itself already represents a significant reduction and we could ask with the same justification whether human expertise which is based primarily on skill can be described adequately using natural language.

Wing (1990) argue persuasively against this myth in the following way: “Programs, however, are formal objects, susceptible to formal methods. . . Thus, programmers cannot escape from formal methods. The question is whether they work with informal requirements and formal programs, or whether they use additional formalism to assist them during requirements specification.”

2.2.2 *Myth 2: Formal languages are difficult to learn*

This is the usual argument of the assembler programmer against FORTRAN and fundamentally boils down to the question of what is easier to understand: low-level or high-level programming languages. Every executable, and in the strictest sense, also every formal specification language can be considered as a programming language at a very high conceptual level. The aspect of an implementation which is efficient and close to the machine is here less important than the interest in offering language constructs at a high conceptual level. Zave (1991), who was deeply involved in the development of the operational specification language PAISley, referred to the fact that: “an executable specification language is a specialized programming language.” The relationship between specification languages and programming languages has changed continually. Or, formulated differently: the current specification languages are the programming languages of the future (regardless of their executability), as the higher conceptual level allows programming which is more efficient because it is more understandable.

The above argument holds for SE in general, but KE itself is in an even stronger position to dispel this myth. The formal languages developed for KE use the model of expertise as a conceptual model. They allow a graphic representation of most modelling primitives and the formal specification can tie onto already existing semi-formal models and refine them. The graphic modelling primitives of the KADS model are supplemented in this, for example, by Petri networks or extended entity relationship diagrams (EER).

2.2.3 *Myth 3: Formal specifications are too complex and difficult to understand*

Basically, the aspects of precision and clarity are discussed here. Because formal and executable specifications force a detailed formulation of the modeled knowledge, formal specifications soon become very large and confusing. Structuring, hierarchization, and modularization are well-known techniques for moderating this problem. The conceptual models used by KE languages offer clear advantages, particularly in this point, as they differentiate clearly between different types or levels of knowledge and subdivide the entire specification into small elements with clearly defined purposes (e.g., an elementary inference action, a domain view of an inference action, a task, a knowledge role, etc.). This again is a clear advantage of the KE languages discussed here over the general-purpose specification languages of SE with their weaker conceptual model. We must always emphasize that a formal specification cannot replace an informal specification, but is intended to refine or supplement it. It is almost always easier to understand an idea by reading a natural language text than to attempt to extract it laboriously from lots of formulae. On the other hand, the attempt at finding a precise and complete natural language definition for a complex problem generally leads to illegible and incomprehensible multi-clause sentences which always contain ambiguities, redundancies, breaks, and contradictions.

A further rebuttal of this point resembles the argument against Myth 1: if formal specifications are too complex and difficult to understand, then this must surely apply with equal if not more force to the lower level implementation languages used in everyday practice, since the programs expressed in these languages are often orders of magnitude larger than the corresponding specification in a high level formal language.

2.2.4 *Myth 4: Formal methods are all about proving programs correct*

The point is made in Hall (1990) that the value of formal specifications is often not in the final product (the actual specification text), but rather in the process of creating that product. This process, with its enforced precision and detail, reveals many insights in the nature of the system to be produced, and these insights, rather than the actual final document, justify the formal

specification effort. In industrial practice, it turns out that the benefits obtained during the process of formal specification are already so crucial that large specification projects report the successful use of formal methods without ever doing a single line of proof (again, see Hall, 1990). It is important not to see specifications only as a useful product, but to see the writing of specifications as an important process. In this process, statements are made more precise, ambiguities and contradictions are recognised and eliminated. The errors and problems which are recognised can be used to improve the informal specifications. Thus occasionally a formal specification is not necessary as a final document, but only the process of formalisation and the improvement of the informal specification which it enables. Wing says on this subject: "The greatest benefit in applying a formal method often comes from the process of formalising it rather than from the end result" (Wing, 1990).

2.2.5 *Myth 5: Formal methods are too expensive and time consuming to use*

In fact the formalisation of a specification does generate additional costs. On the other hand, this objection is probably the most superficial. Every standard textbook of SE is full of admonitions that the later a specification or programming mistake is discovered in the software development process, the more expensive it will be to eliminate it. An informal specification with its basic incompleteness forces the programmer literally to bridge gaps using his imagination. And he or she does not always meet the expectations of the later users or the lacking expert knowledge.

The predominant share of the effort in the construction of a formal specification is thus not due to the use of a formal specification, but rather to the goal of making an informal specification more precise and eliminating ambiguities and contradictions. If the system is to be implemented, then this effort must be expended anyway. But this will then no longer be carried out in an appropriate phase and by the appropriate persons. "The fact is that writing a formal specification decreases the cost of development" (Hall, 1990).

Besides these general points, again KE is in a particularly strong position to dispel this myth. As shown in van Harmelen & Aben (1995), the close correspondence between the informal conceptual model and the formal specification language in KE can be exploited to give a much stronger support (both automated and non-automated) to the transition from informal to formal model, and such support is likely to significantly reduce the cost associated with formal specification.

Finally, the KE community has traditionally placed great emphasis on the possibility to *re-use* library elements or fragments of existing models in the construction of new models. Such re-use is only properly possible if the fragments are indeed specified precisely enough to be reusable. In the context of knowledge engineering, these fragments will consist of the structure (ontology) of domain-knowledge, elementary inference steps (Aben, 1995) or entire problem-solving methods (Breuker & Van de Velde, 1994). Each of these fragments must be accompanied by a precise characterisation of the assumptions it makes about the knowledge, data and computational environment required for its proper functioning. Re-use of such fragments will then (among other things) require a (preferably formal) proof that the requirements of a particular fragment are indeed fulfilled in the given circumstances. Such an increased possibility of reuse will be a great pay-back for the additional effort involved in building the final specification of the building blocks.

2.2.6 *Myth 6: Formal methods never worked in SE, so why would they work in KE?*

First of all, the premise of this myth is increasingly untrue as time proceeds. Highly publicized accounts of the application of formal methods to a number of well-known systems, e.g., the Darlington Nuclear Facility in the UK (Anderson & Bruns, 1995; Parnas, 1995), the CICS project at IBM-UK (Hoare, 1995) and Airbus (Peleska et al., 1995), all reported in Hinchey and Bowen (1995). Craigen et al. (1993) and Craigen (1995) provide a useful survey of the use of formal methods in SE.

Secondly, as already stated above, the formal languages recently developed for KE purposes are tightly coupled with the underlying conceptual models developed earlier in the field. This has as

advantages (i) that these languages can give strong support for the construction of a formal specification, and (ii) that the structural correspondence can be used to verify that the formal specification does indeed capture the informally stated requirements.

The same argument about the close correspondence between informal and formal models can also be used to rule out the following:

2.2.7 *Myth 7: Formal methods are unacceptable to users*

In particular the graphical representations available for many formal languages in KE (see, for example, MIKE (Angele et al., 1993)) and the possibility to interpret the formal constructions in terms of the underlying conceptual model form a significant bridge to users and domain experts alike.

2.2.8 *Myth 8: Formal methods are only useful in safety-critical systems*

It is true that safety-critical applications can gain obvious benefits from formal methods. We quote the following from Bowen and Hinchey (1994):

The UK Ministry of Defence (MoD) draft Interim Defence Standards 00-55 and 00-56 mandate the extensive use of formal methods. Standard 00-55 sets forth guidelines and requirements; the requirements include the use of a formal notation in the specification of safety-critical components, and an analysis of such components for consistency and completeness. All safety-critical software* must also be validated and verified; this includes formal proof and rigorous (but informal) correctness proofs, as well as more conventional static and dynamic analysis. Standard 00-56 deals with the classification and hazard analysis of the software and electronic components of defence equipment, and also mandates the use of formal methods.

The Atomic Energy Control Board (AECB) in Canada, in conjunction with David Parnas at McMaster University, has commissioned a proposed standard for software for computers in the safety systems of nuclear power stations. Ontario-Hydro has developed a number of standards and procedures within the framework set by AECB and further procedures are under development.

Many reported applications concern safety-critical systems; for example, Hinchey and Bowen (1995) mention railway signalling and railway traffic tracking, Airbus cabin communication systems, besides some of the applications mentioned above. Nevertheless, non-safety-critical applications have been reported as well: Hinchey and Bowen (1995) mention instrumentation systems, telephone switching systems, secure operating systems and microprocessors. Famous examples of these include the T800 Transputer floating-point unit and parts of the T9000 transputer pipeline architecture (all of these reported in Bowen & Hinchey, 1994). These real-life examples also dispel the following myth.

2.2.9 *Myth 9: Formal methods are not used on real large-scale systems*

Besides the extensive lists of applications mentioned above, it is important to realise here that formal methods need not necessarily be applied to entire systems. It is often possible to isolate crucial, complicated or critical components of a large system, and to limit the use of formal methods to these subsystems.

We must admit at this point, however, that the applications listed above concern formal methods from SE. Formal methods in KE have not yet reached the degree of maturity that they have made their way into industrial applications, and this must be seen as one of the major challenges of the field in the years to come.

3 Current state-of-the-art

The work on formal methods in KE has been recently and extensively surveyed in Treur (1994) and Fensel and van Harmelen (1994), and we refer the interested reader to these works for detailed comparisons between the different approaches. In this section, we will instead give a broader, less

*Safety critical software is defined as software which *on its own* could cause a significant accident.

Table 1 A brief survey of KADS-based formal languages

<i>Language</i>	<i>Site</i>	<i>Language constructs</i>	<i>Executable</i>	<i>Static semantics</i>	<i>Dynamic semantics</i>	<i>Active*</i>
OMOS (Linster, 1993)	GMD, Bonn	frames, hierarchies predicates, functions	yes	no	no	no
MODEL-K (Karch et al., 1993)	GMD, Bonn	frames, hierarchies predicates, functions	yes	no	no	no
MoMo (Voss & Voss, 1993; Gebhardt, 1995)	GMD, Bonn	Petri nets, functions, arbitrary KR-languages	yes	no	no	yes
FORKADS (Wetter & Schmidt, 1991; Wetter, 1990)	IBM, Heidelberg	sorted logic	yes	Tarskian models	no	no
KARL (Angele et al., 1995; Fensel, 1995)	Univ. of Karlsruhe	F-logic Horn logic Dynamic logic	yes	perfect Herbrand models	Kripke models	yes
(ML) ² (van Harmelen & Balder, 1992)	Univ. of Amsterdam	sorted logic meta-logic dynamic logic	no	Tarskian models	Kripke models	yes
QIL (Shadbolt et al., 1992)	Univ. of Nottingham	predicate logic temporal logic epistemic logic	no	Tarskian models	Kripke models	yes
K _{BS} SF (Veld et al., 1994)	PTT Labs, Netherlands	algebra, sorted logic procedures	no	Tarskian models, initial algebras	Plotkin style	no
TFL (Pierrot-Goldbreich & Talon, 1995)	Univ. de Paris Sud	algebraic data types	no	loose algebraic semantics	loose algebraic semantics	yes
GCLA (Aronsson et al., 1994)	SICS, Sweden	generalised logic programs	yes	partial inductive definitions	partial inductive definitions	yes

*This column indicates whether the language is still the subject of active research, development and use, or whether work on the language has stopped.

detailed and also somewhat more recent overview of the current state of the field. Based on this overview, in the final section of this paper we will identify the most important and urgent points on the agenda of this community.

The work in the field until now is best summarised by the catalogue of formal languages that have been developed in recent years. Three groups of work can be distinguished here: (i) formal languages based on the KADS conceptual model; (ii) formal languages based on other conceptual models; and (iii) the use of formal languages from software engineering.

The languages of group (i) have been surveyed and compared in detail in Fensel and van Harmelen (1994). That paper compares languages by looking at their operational aspects, their epistemological commitments and their formal underpinning. A very brief survey of that paper is given in Table 1. Analysis in that paper shows that a major determinant of choices made in the




Figure 1 Flexibility of reasoning patterns

various languages was whether a language was intended to be operational or not. Some languages (OMOS, MODEL-K, MoMo) aim at clarifying the meaning of conceptual models through operationalising such models, and do not aim at providing a formal semantics. Other languages ((ML)², QIL, K_{BS}SF) aim at formalising conceptual models, but use such expressive formal constructions that this precludes effective executability. This non-executability is not only a matter of efficiency, but often also of the much more principled issue of (semi)-decidability. A third group of languages (KARL, FORKADS) aims to combine executability and formalisation. This main aim of language was found to explain many of the differences found in the detailed comparisons between the languages presented in Fensel and van Harmelen (1994).

This detail in the comparison was possible because all the languages discussed in Fensel and van Harmelen (1994) were based on the same underlying conceptual model. This was not the case for the languages compared in Treur & Wetter (1993). The basis of that comparison was a specific task-and domain-model (a simple timetable scheduling task) which was modelled in a number of languages. Again, the purpose of the language (operationalisation, formalisation or a combination of these two) was found to be of prime importance. The following common properties were found to hold across the set of languages discussed in Treur & Wetter (1993):

- composing a complex specification out of components which are each declaratively specified;
- distinctions between static and dynamic aspects of a specification;
- distinctions between generic and domain-specific parts of the specification.

A coarse comparison was done among the languages, based on three dimensions:

- **Expressive power for domain knowledge:** The main distinction here was between languages which used a sublanguage of first-order predicate logic (such as Horn logic) and those which used full first-order logic. Other languages which were used (such as modal logic) could in principle be encoded within first-order order logic so that from a strict semantic point of view, these languages are equivalent.
- **Flexibility of reasoning patterns:** This property is summarised in Figure 1. It indicates that AIDE has a completely fixed structure of reasoning pattern, (ML)² and KARL have a fixed overall structure but can be configured within this structure: the three-layer structure of these languages (based on KADS) is fixed, but they are configurable within each of these three layers. MC and DESIRE are fully configurable and impose no fixed configuration.
- **Expressiveness of control knowledge:** Treur & Wetter (1993) agreed with Fensel and van Harmelen (1994) that this is the main point of difference between the various languages, and the extent to which the languages succeeded in formalising the control knowledge is reported in Fensel and van Harmelen (1994) and shown here in Table 2.

A more detailed comparison was also carried out in Fensel and van Harmelen (1994) and is summarised in Table 2.

The two survey papers discussed above both focused on novel specification languages specifically designed for use in knowledge engineering. Another (although much smaller) body of work

Table 2 Detailed comparison as in Fensel and van Harmelen (1994)

	(ML) ² (van Harmelen & Balder, 1992)	MC (Giunchiglia & Traverso, 1995)	AIDE (Kassel & Greboval, 1993)	KARL (Angele et al., 1995; Fensel, 1995)	DESIRE (van Langevelde et al., 1992)	OBJ3 (Nakagawa et al., 1993)	MILORD (Sierra & Godo, 1993)	K _{BS} SF (Veld et al., 1994)
1	FOL meta-logic dyn. logic	FOL meta-logic	restricted FOL	Horn logic, dyn. logic	3-valued FOL meta-logic temp. logic	order sorted algebra	multi-valued logic locally	order-sorted logic only at domain level
2	yes	yes	only at domain level	yes	component-wise. temp. composition	only at domain level		
3	yes	by the user	yes	yes	yes	no	yes	yes
4	yes	yes	yes	yes	yes	yes	yes	yes
6	yes	yes	no	yes	yes	yes	yes	hard
7	no	yes	no	no	yes	no	yes	via paramaterised deduction
8	partial	no	limited	yes	partial	yes	yes	partial
9	yes	yes	partial	yes	yes	yes	yes	yes

1. Expressive power
2. Declarativeness
3. Adequacy to specify dynamic aspects of reasoning patterns
4. Possibility to specify multi-level architecture
6. Adequacy to specify non-classical reasoning
7. Possibility to specify integrated systems
8. Executability
9. Availability of formal semantics

has concentrated on using existing languages from software engineering to problems in knowledge engineering. Examples of this approach are:

- Craig (1991a,b): specifications of a blackboard architecture (CASSANDRA) and a production-rule architecture (including meta-rules, ELEKTRA), using Z (Spivey, 1992, 1988).
- Velthuisen (1992); Velthuysen & Braspenning (1991): specifications of a large number of blackboard architectures using CCS (Milner, 1989).
- Nakagawa et al. (1993): specification of a simple scheduling task (from Treur, 1994) using OBJ3 (Futatsugi et al., 1985).
- Todd & Stamper (1993): specification of diagnostic reasoning patterns using Z.
- Krause et al. (1993): specification of medical knowledge models using Z.
- Milnes (1992): specification of the soar cognitive architecture in Z.

4 Future research directions

As the concluding section of this survey paper, we will sketch what we see as the most important and/or challenging research problems for this community in the next few years.

4.1 Formalisms for dynamics

As we have seen above, both survey papers (Treur & Wetter, 1993; Fensel & van Harmelen, 1994) identified formalisms for dynamics as the cause of the largest differences between the currently available specification languages. In software engineering and information systems design we also see a wide variety of formalisms to deal with dynamic behaviour of systems (e.g., evolving algebras (Gurevich, 1993), dynamic logic (Harel, 1984) and variants thereof (Spruit et al., 1995), Petri nets (Jensen, 1987), temporal logics (Treur, 1994)). These are discussed in some detail in the accompanying paper by Fensel (1995). It is at the moment unclear what the strengths and

weaknesses of the various formalisms are, and what the actual requirements of knowledge engineering specification languages are in this respect.

4.2 Semantics

Connected with the open question on appropriate formalisms for dynamic behaviour of KBS is the semantics provided for such formalisms. Again, the set of currently available languages differ widely in this respect, and even more options are available in the literature on software engineering and on logic.

A second problem with the semantics of these languages is the following: almost all languages view a complex specification as built out of simpler components in a structured way. Good traditional semantic foundations exist for most of the individual components employed in these languages, but it is often unclear how the overall semantics should be composed out of these individual components. Foundational work is required for a clear view on such compositional semantics.

4.3 Proof calculi

Clearly related to the problem with semantics is the development of proof calculi for these languages. Notwithstanding our reply to myth 4 above, in order to reap the full benefits of these formal specification languages, we need proof calculi which allow proving general properties of specifications. In general, a proof calculus is required which enables us to prove properties of the form.

$$\text{preconditions} \wedge \text{specification} \rightarrow \text{post-conditions}$$

where *preconditions* and *post-conditions* are sets of first-order formula, and *specification* is a formal specification which includes static as well as dynamic aspects of a system. That is, *specification* is a set of formulae in a logic which includes states and changes of states. Examples of this type of logic are dynamic and temporal logic. A formal reasoning calculus together with automated proof support elicits the full power of formalisation as it becomes possible to prove implicit properties of a specified system (e.g., Heisel et al., 1988, 1990). Approaches in software engineering like VDM already deliver good results in this area (Bicarregui et al., 1994).

4.4 Validation and verification

Closely related to such proofs of properties of a specification is the use of these specification languages for validation and verification. This defines a strong link to work done in program verification (see Colburn et al., 1993; Fetzer, 1988 for a survey and Heisel et al., 1990, for an approach using theorem proving in dynamic logic), and to work done in validating and verifying KBS (cf. Ayel & Laurent, 1991; Lydiard, 1992). Until now, most of the work in validating and verifying knowledge-based systems has focused on specific implementation formalisms, but this could change now as these techniques can be applied to conceptual models when these models are supplemented by a formal language and semantics. Furthermore, the current techniques only apply to the static semantics of knowledge, and do not take into account the control structures under which this knowledge is to be used. The paper by Meseguer and Preece (1995) in this issue provides a discussion on the relationship between conceptual modelling, formal methods and validation and verification of KBS.

4.5 Interaction with environment

A weak aspect of many if not all of the current knowledge engineering specification languages is a limitation already inherent in the conceptual models on which these languages are based. This

concerns the modelling of the interaction of a system with its environment. Issues such as interaction with a user, interactions with asynchronous events and interaction with other software systems in a heterogeneous environment have all remained very underdeveloped until now. The task- and cooperation model in KADS (Breuker & de Greef, 1993) is a step in this direction, but not nearly as fully developed as the corresponding conceptual model of expertise. Work on DESIRE also has taken the interaction with the environment into account and is ahead of KADS in this respect.

This is all the more urgent since each of these points is likely to be a significant factor in many application areas. Extensive work remains to be done here, both at the conceptual and at the formal level.

4.6 Reusable building blocks

Reuse of software is a means to improve both its quality and the efficiency of its development process. In the context of knowledge engineering, reuse is achieved by domain ontologies, libraries of problem-solving methods (Breuker and Van de Velde, 1994), and libraries of elementary inference actions (Aben, 1995). Compared to the effort which was put in formal specification techniques for conceptual models of KBS much less effort has been spent to develop formalisms and methods for formal support in developing and using reusable components. An exception is the work of Aben (1995), who provides a formal framework for reusing elementary inference actions.

Fensel (1994b) provides an informal analysis of a formally specified problem-solving method taking into account its reusability. The main results are: First, problem-solving methods make assumptions about the domain knowledge which are required by these methods to solve a given problem effectively and efficiently. Therefore, several types of proof obligations arise in the context of reusing such building blocks. Second, reusable domain theories have to be supplemented by meta-layer descriptions to make properties of the specified domain knowledge explicit. This enables the proof whether a domain knowledge fulfills the assumptions of a problem-solving method. Third, a problem-solving method defines in fact a whole family of different variants with different assumptions, elementary inference steps, and overall control. Therefore, a library can only contain general schemas and a refinement calculus is required to support the development of an appropriate instantiation of such a method for a given application. Fourth, appropriate specification formalisms are required which reflect the needs of describing reusable building blocks.

4.7 Exposure to applications

In our replies to myths 6 and 8, we were forced to use almost exclusively examples from applications of formal methods in software engineering rather than knowledge engineering. The reason for this is of course that the use of formal methods in knowledge engineering has simply not yet reached the stage of industrial application. Of all the open points of future work in this section, this point is perhaps the most important one. After the development of a variety of formal specification languages, many of which now seem to be stable, the time is ripe for proving the value of these languages in real-life applications. Areas where KBS are already used and which seem appropriate for the use of these formal techniques are the process-industry (e.g., chemical plants, energy production), health (e.g., medical decision-making (Fox, 1993)), transport and finance.

Acknowledgements

We are grateful to Jan Treur, Ameen Abu-Hanna, Pascal van Eck and two anonymous referees for their comments on this paper.

References

- Aben, M, 1995, *Formal Methods in Knowledge Engineering* PhD thesis, University of Amsterdam, Faculty of Psychology, February.
- Anderson, S and Bruns, G, 1995, "The formalization and analysis of communication protocol" In: M Hinchey and J Bowen (eds.), *Applications of Formal Methods*, Prentice Hall.
- Angele, J, Fensel, D, Landes, D, Neubert S and Studer, R, 1993, "Model-based and incremental knowledge of engineering: The MIKE approach" In: J Cuenca (ed.), *Knowledge Oriented Software Design*, Number A-27, IFIP Transactions, North Holland.
- Angele, J, Fensel, D, Landes, D and Studer, R, 1995, "The knowledge-based acquisition and representation language KARL" (submitted for publication).
- Aronsson, M, Eriksson, M and Kreuger, P, 1994, "GCLA and a sketch of its use for representing KADS models" In: *Proceedings of the ECAI'94 Workshop on Formal Specification Methods for Knowledge-Based Systems* Amsterdam, August.
- Ayel, M and Laurent, J-P, editors, 1991, *Validation, Verification and Test of Knowledge-Based System* Wiley.
- Barbuceanu, M, 1993, "Towards integrated knowledge modeling environments" *Knowledge Acquisition* 5 (3).
- Bauer, C and Karbach, W, editors, 1992, *Proceedings Second KADS User Meeting* ZFE BT SE 21, Otto-Hahn Ring 6, D-8000 Munich 83, 17-18 February. Siemens AG.
- Bicarregui, J, Fitzgerald, J, Lindsay, P, Moore, R and Ritchie, B, 1994, *Proof in VDM: A Practitioner's Guide* Springer-Verlag.
- Bowen, J and Hinchey, M, 1994, "Seven more myths of formal methods" In: M Naftalin, T Denvir and M Bertran (eds.), *Proceedings of the 2nd International Symposium of Formal Methods Europe (FME'94)* volume 873 of *Lecture Notes in Computer Science* pp 105-115, Barcelona, October.
- Breuker, JA and de Greef, P, 1993, "Modelling system-user cooperation in KADS In: ATh Schreiber, BJ Wielinga and JA Breuker (eds.), *KADS: A Principled Approach to Knowledge-Based System Development* pp 47-70, Academic Press.
- Breuker JA and Van de Velde, W, editors, 1994, *The Common KADS Library for Expertise Modelling* IOS Press.
- Breuker, JA, Wielinga, BJ, van Someren, M, de Hoog, R, Schreiber, ATh, de Greef, P, Bredeweg, B, Wielemaker, J, Billault, JP, Davoodi, M and Hayward, SA, 1987, "Model Driven Knowledge Acquisition: Interpretation Models" ESPRIT Project P1098 Deliverable D1 (task A1), University of Amsterdam and STL Ltd.
- Buchanan, BG and Shortliffe, EH, 1984, *Rulebased Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project* Addison-Wesley.
- Chandrasekaran, B, 1986, "Generic tasks in knowledge based reasoning: High level building blocks for expert system design" *IEEE Expert* 1 (3) 23-30.
- Clancey, WJ, 1993, "Heuristic classification" *Artificial Intelligence* 27 289-350.
- Colburn, TR, Fetzer, JH and Rankin, TL, editors, 1993, *Program Verification* Kluwer Academic.
- Craig, I, 1991a, *The Formal Specification of Advanced AI Architecture: Lecture Notes in Mathematics* Ellis Horwood.
- Craig, I, 1991b, "The formal specification of ELEKTRA" Technical report, Department of CS, University of Warwick.
- Craigen, C, 1995, "Applications of formal methods: Observations and trends" In: M Hinchey and J Bowen (eds.), *Applications of Formal Methods* Prentice Hall.
- Craigen, D, Gerhart, S and Ralston, T, 1993, "An international survey of industrial applications of formal methods" Technical report, U.S. Department of Commerce, Technology administration, National Institute of Standards and Technology, Computer Systems Laboratory, Gaithersburg, MD 20899, USA, March.
- Fensel, D, 1994, "Sinn und unsinn formaler spezificationsprachen fuer wissensbasierte systeme" *Kuenstliche Intelligenz* (4) (in German).
- Fensel, D, 1994b, "A case study on assumptions and limitations of a problem solving method", *9th Banff Knowledge Acquisition for KBS Workshop*, Banff, Canada, February.
- Fensel, D, 1995, *The Knowledge-Based Acquisition and Representation Language KARL*, Kluwer.
- Fensel, D and van Harmelen, F, 1994, "A comparison of languages which operationalise and formalise KADS models of expertise" *The Knowledge Engineering Review* 9 105-146.
- Fetzer, JH, 1988, "Program verification: The very idea" *Communications of the ACM* 31 September.
- Fox, J, 1993, "On the soundness and safety of expert systems" *Artificial Intelligence in Medicine* 5 159-179.
- Futatsugi, K, Goguen, J, Jouannaud, J and Meseguer, J, 1985, "Principles of OBJ2" In: BK Reid (ed.) *Proceedings of Twelfth Symposium on Principles of Programming Languages* pp 52-66. ACM.
- Gebhardt, F, 1995, "MoMo: language and case studies" Fabel-Report 36, GMC, Sankt Augustin, July.

- Giunchiglia, E and Traverso, P, 1995, "A multi-context architecture for formalizing complex reasoning" *International Journal of Intelligent Systems* 10 (5) 501–539, May.
- Giunchiglia, E, Traverso, P and Giunchiglia, F, 1993, "Multi-context systems as a specification framework for complex reasoning systems" In J Treur and Th Wetter (eds.), *Formal Specification of Complex Reasoning Systems* Ellis Horwood.
- Gurevich, Y, 1993, "Evolving algebras, a tutorial introduction" In EG Rozenberg et al. (eds.), *Current Trends in Computer Science* World Scientific.
- Hall, JA, 1990, "Seven myths of formal methods" *IEEE Software* 7 (5) 11–19, September.
- Harel, D, 1984, "Dynamic logic" In D Gabbay and F Guenther (eds.), *Handbook of Philosophical Logic, Vol. II: extensions of Classical Logic* pp 497–604. Reidel.
- Heisel, M, Reif, W and Stephan, W, 1988, "Implementing verification strategies in the KIV-system" In R Lusk and R Overbeek (eds.), *9th International Conference on Automated Deduction (CADE'88)* Number 310 in *Lecture Notes in Computer Science* pp 131–140, Argonne, Springer-Verlag.
- Heisel, M, Reif, W and Stephan, W, 1990, "Tactical theorem proving in program verification" In: *Proceedings of the 10th International Conference on Automated Deduction (CADE'90)*, Volume 449 of *Lecture Notes on Computer Science* Springer-Verlag.
- Hinchey, M and Bowen, J, editors, 1995, *Applications of Formal Methods* International Series in Computer Science. Prentice-Hall.
- Hoare, J, 1995, "Formal development of CICS with B" In: M Hinchey and J Bowen (eds.), *Applications of Formal Methods* Prentice-Hall.
- in't Veld, L, Jonker, W and Spee, JW, 1994, "The specification of complex reasoning tasks in $K_{BS}SF$ " In J Treur and Th Wetter (eds.), *Formal Specification of Computer Reasoning Systems* pp 233–256 Ellis Horwood.
- Jackson, P, 1990, *Introduction to Expert Systems* Addison Wesley.
- Jaurent, JP, Ayel, J, Thome, F and Ziebelin, D, 1986, "Comparative evaluation of three expert system development tools: KEE, knowledge craft, ART" *Knowledge Engineering Review* 1 (4) 19–29.
- Jensen, K, 1987, "Coloured petri nets" In W Brauer, W Reisig and G Rozenberg (eds.), *Applications and Relationships to Other Models of Concurrency, Advances in Petri Nets 1986 Part I*, volume 254 of *Lecture Notes of Computer Science* pp 248–299. Springer-Verlag.
- Jonker, W, Spee, JW, in't Veld, L and Koopman, M, 1991, "Formal approaches towards design in SE and their role in KBS design" In: *Proceedings of the IJCAI'91 Workshop on Software Engineering for Knowledge-Based Systems* Sydney, Australia, August.
- Karbach, W, Voß, A, Schuke, R and Drouwen, U, 1991, "Model-K: Prototyping at the knowledge level" In: *Proceedings Expert Systems-91* pp 501–512, Avignon, France.
- Kassel, G and Greboval, C, 1993, "How AIDE succeeds in an example design task" In H Treur and Th Wetter (eds.), *Formal Specification of Complex Reasoning Systems* Ellis Horwood.
- Keene, SE, 1984, *Object-Oriented Programming in Common Lisp* Addison-Wesley.
- Krause, P, Fox, J, O'Neill, M and Glowinski, A, 1993, "Can we formally specify a medical decision support system" *IEEE Expert* 8 (3) 56–61.
- Landes, D and Studer, R, 1995, "The treatment of non-functional requirements in MIKE" In: *Proceedings of the 5th European Software Engineering Conference (ESEC'95)* Barcelona, Spain.
- Linster M, 1993, "Using the operational modeling language OMOS to represent KADS conceptual models" *Knowledge Acquisition*.
- Löckenhoff, C, Fensel, D and Studer, R, editors, 1993, *Proceedings of the Third KADS User Meeting ZFE BT SE 21*, Otto-Hahn Ring 6, D-8000 Munich 83, 8–9 March Siemens AG.
- Lydiard, T, 1992, "Overview of current practice and research initiatives of the verification and validation of KBS" *Knowledge Engineering Review* 7 (2) 101–113.
- Marcus, S, editor, 1988, *Automatic Knowledge Acquisition for Expert Systems* Kluwer.
- Milner, R, 1989, *Communication and Concurrency* International Series in Computer Science, Prentice-Hall.
- Milnes, B, 1992, "A specification of the Soar cognitive architecture in Z", Technical Report CMU-CS-92-169. Carnegie Mellon University, August.
- Musen, MA, 1989, *Automated Generation of Model-Based Knowledge-Acquisition Tools* Pitman.
- Nakagawa, A, Sakakihara, T and Futatsugi, K, 1993, "Algebraic specification of reasoning systems" In: J Treur and Th Wetter (eds.), *Formal Specification of Complex Reasoning Systems* Ellis Horwood.
- Newell, A, 1982, "The knowledge level" *Artificial Intelligence* 18 87–127.
- Parnas, D, 1995, "Using mathematical descriptions in the inspection of safety-critical software" In: M Hinchey and J Bowen (eds.), *Applications of Formal Methods* Prentice-Hall.
- Peleska, J, Hamer, U and Hoercher, H-M, 1995, "The airbus a330/340 cabin communication system—a Z application" In: M Hinchey and J Bowen (eds.), *Applications of Formal Methods* Prentice-Hall.
- Pierret-Goldbreich, C and Talon, X, 1995, "An algebraic specification of the dynamic behaviour of knowledge-based systems" *The Knowledge Engineering Review* (submitted).

- Shadbolt, N, Aitken, S and Reichgelt, LH, 1992, "Representing kads models in QIL" Working Paper WP-006, AI Group, University of Nottingham.
- Schreiber, Ath and Terpstra, P, 1995, "Sisyphus-VT: A Common KADS solution" *International Journal of Human-Computer Studies* (in press).
- Sierra, A and Godo, L, 1993, "Specifying simple scheduling tasks in a reflective and modular architecture" In: J Treur and Th Wetter, editors, *Formal Specification of Complex Reasoning Systems* Ellis Horwood.
- Spivey, JM, 1988, *Understanding Z: A Specification Language and its Formal Semantics* volume 3 of *Cambridge Tracts in Theoretical Computer Science* Cambridge University Press.
- Spivey, JM, 1992, *The Z Notation: A Reference Manual* Series in Computer Science, Prentice Hall.
- Spruit, P, Wieringa, R and Meyer, J-J, 1995, "Axiomatization, declarative semantics and operational semantics of passive and active updates in logic databases" *Journal of Logic and Computation* 5 (1).
- Todd, B and Stamper, R, 1993, The formal design and evaluation of a variety of medical diagnostic programs" Technical Monograph PRG-109, Oxford University Computing Laboratory, September.
- Treur, J, 1994, "Temporal semantics of meta-level architectures for dynamic control" In: L Tribourg and F Turini (eds.), *Logic Program Synthesis and Transformation—Meta-Programming in Logic (LOPSTR'94-META'94)* volume 883 of *Lecture Notes in Computer Science* pp 353–376, Springer-Verlag.
- Treur, J and Wetter, Th, editors, 1993, *Formal Specification of Complex Reasoning Systems* Workshop Series, Ellis Horwood.
- Treur, J and Willems, M, 1994, "A logical foundation for verification" In T Cohn (ed.), *Proceedings of the 11th European Conference on AI (ECAI'94)* pp 745–749, August. Wiley.
- Ueberreiter, B and Voß, A, editors, 1991, *Materials KADS User Meeting, February 14/15 1991* Siemens AG ZFE IS INF 32, Munich Perlach, Germany (in German).
- van Harmelen, F and Aben, M, 1995, "Structure preserving specification languages for knowledge-based systems" *International Journal of Human Computer Studies* (formerly *Journal of Man Machine Studies*).
- van Harmelen, F and Balder, JR, 1992, "(ML)²: a formal language for KADS models of expertise" *Knowledge Acquisition* 4 (1). (Special issue: 'The KADS approach to knowledge engineering', reprinted in *KADS: A Principled Approach to Knowledge-Based System Development*, 1993, Schreiber, Ath, et al., editors).
- van Langevelde, LA, Philipsen, AW and Treur, J, 1992, "Formal specification of compositional architecture" In: B Neumann (ed.), *Proceedings ECAI'92* Vienna, pp 272–276. Wiley. (Longer version available as: Report IR-282, Mathematics and Computer Science, Free University of Amsterdam.)
- van Melle, W, 1979, "A domain independent production rule system for consultation programs" In: *IJCAI-79* pp 923–925.
- Velthuisen, H, 1992, "The nature and applicability of the blackboard architecture" PhD thesis, University of Limburgh, Maastricht, The Netherlands.
- Velthuisen, H and Braspenning, P, 1991, "A conceptual and formal study of the blackboard architecture" In: E Moskilde (ed.), *Proceedings of the European Simulation Multi-Conference (ESM'91)* pp 374–379, Copenhagen. The Society for Computer Simulation.
- Voss, H and Voss, A, 1993, "Reuse-oriented knowledge engineering with MoMo" In: *Proceedings of the 5th International Conference on Software Engineering and Knowledge Engineering (SEKE'93)* San Francisco, June.
- Wetter, J, 1990, "First-order logic foundation of the KADS conceptual model" In: BJ Wielinga, J Boose, B Gaines, G Schreiber and M van Someren (eds.), *Current Trends in Knowledge Acquisition* pp 356–375, Amsterdam, The Netherlands, May. IOS Press.
- Wetter, Th and Schmidt, W, 1991, "Formalization of KADS interpretation models" In: L Steels and B Smith, editors, *AISB91, Artificial Intelligence and Simulation of Behavior* Springer-Verlag.
- Wielinga, BJ and Breuker, JA, 1986, "Models of expertise" In: *Proceedings ECAI-86*, pp 306–318.
- Wielinga, BJ, Schreiber, Ath and Breuker, JA, 1992. "KADS: A modelling approach to knowledge engineering" *Knowledge Acquisition* 4 (1) 5–53. (Special issue 'The KADS approach to knowledge engineering'. Reprinted in: B Buchanan and D Wilkins, editors, 1992, *Readings in Knowledge Acquisition and Learning* Morgan Kaufmann, pp 92–116.)
- Wing, JM, 1990, "A specifier's introduction to formal methods" *IEEE Computer* 23 (9) 10–24, September.
- Zave, P, 1991, "An insiders evaluation of PAISLey" *IEEE Transactions on Software Engineering* 18 (3) 212–225.