

Assessing the Role of Formal Specifications in Verification and Validation of Knowledge-Based Systems

Pedro Meseguer

IIIA - Artificial Intelligence Research Institute

CSIC - Spanish Scientific Research Council

Campus Universitat Autònoma de Barcelona

08193 Bellaterra, SPAIN

Phone +34-3-580 95 70; FAX +34-3-580 96 61

Email: meseguer@lsi.upc.es

Alun D. Preece

University of Aberdeen

Computing Science Department

Aberdeen AB9 2UE, Scotland, UK

Phone +44 1224 272296; FAX: +44 1224 273422

Email: apreece@csd.abdn.ac.uk

Abstract

This paper examines how formal specification techniques can support the verification and validation (V&V) of knowledge-based systems. Formal specification techniques provide levels of description which support both verification and validation, and V&V techniques feed back to assist the development of the specifications. Developing a formal specification for a system requires the prior construction of a conceptual model for the intended system. Many elements of this conceptual model can be effectively used to support V&V. Using these elements, the V&V process becomes deeper and more elaborate and it produces results of a better quality compared with the V&V activities which can be performed on systems developed without conceptual models. However, we note that there are concerns in using formal specification techniques for V&V, not least being the effort involved in creating the specifications.

Keywords

Knowledge-based systems, artificial intelligence, formal specification, verification, validation, life-cycle.

1 INTRODUCTION

Unlike most conventional software, knowledge-based systems (KBS) are rarely validated and verified against an explicit specification of user requirements. Probably the main reason for this is that the tasks which these systems are required to perform are not well-understood at the outset of development, and so development proceeds in an exploratory manner — usually via

prototyping — until a system is produced which is deemed to embody the implicit requirements of the prospective users. This explains why most of the early efforts to validate KBS by comparing their performance directly against the performance of human performers, rather than against any specification document (Buchanan and Shortliffe, 1984).

The problems with this approach are obvious: the validation is inherently prone to bias and, even when an "acceptable" system is deemed to have been produced, it is not clear what the system actually does (or may fail to do). Much of the work done in recent years to improve the state of verification and validation (V&V) practice for KBS has held on to the assumption that, because KBS are difficult to specify, practical V&V techniques should not depend upon the existence of detailed specification documents. This assumption is clearly seen in, for example, the work done on checking KBS for domain-independent *anomalies* such as inconsistency and incompleteness (Preece, Shinghal and Batarekh, 1992), as well as in the quantitative techniques for comparing KBS with human "experts" (O'Keefe, Balci and Smith, 1987).

Such techniques have been shown to be effective, but to a limited extent only. This paper will examine the limitations of a number of "state-of-the-art" KBS V&V techniques, and will assess how the power of the techniques can be extended when precise specification documents are available for the system. We will focus upon *formal specifications*, because informal and pseudo-formal specifications are too weak to provide a foundation for V&V.*

Therefore, the main aim of this paper is to examine the ways in which formal specification techniques can support the V&V of KBS. In doing so, two related issues naturally come under consideration. Firstly, it becomes apparent that verification and validation techniques can, in turn, support the development of formal specifications. Secondly, it becomes necessary to consider how specification, verification and validation techniques need to be applied within the whole KBS development process. This paper will touch upon these issues, although they will need more detailed consideration in their own right.

Verification and Validation of KBS

Before proceeding, it is necessary briefly to define the terms verification and validation for the purposes of this paper. *Verification* is a process aimed at demonstrating whether a system meets its specified requirements; this is often called "building the system right" (O'Keefe, Balci and Smith, 1987), which we take to mean "checking the system against its (preferably explicit and formal) specifications". *Validation* is a process aimed at demonstrating whether a system meets the user's true requirements — often called "building the right system". Verification and validation can be viewed as a set of techniques and an associated process in which the techniques are applied, as part of the whole development process (for example, static verification of the knowledge base, followed by dynamic testing of the whole KBS (Preece, 1990)).

Formal Specification Techniques for KBS

Like V&V, formal specification techniques for KBS include a number of techniques (formal specification languages of various kinds) and processes (for example, transformation from a pseudo-formal specification to an implementation through several levels of detail). A *formal specification* for a KBS includes (Fensel and van Harmelen, 1994) (i) some specification of the I/O behaviour of the system (establishing the correct relation between data and results) and (ii) a description of how this behaviour can be made. This second element blurs the distinction between specification and design and is perhaps controversial from the standpoint of traditional software specification. However, it is a feature of current KBS formal specification techniques, and can be of great benefit in V&V, so we will not debate its appropriateness here.

*We note, however, that non-formal specifications may be *approximated* by formal specifications and employed for *partial* V&V; this is the role played by the *pseudo-formal* specifications in (Laurent, 1992).

2 VERIFICATION AND VALIDATION OF KBS

The earliest validation technique in AI was Alan Turing's proposal on how to decide if a program could be considered "intelligent", commonly known as the "Turing test" (Turing, 1950). This is a blind test where the evaluator communicates through a teletype with a person and a program; if the evaluator is unable to differentiate between the person and the program, the program is considered to be intelligent. Although many criticisms have been levelled against the Turing test as a general procedure to characterize intelligent behaviour, the idea of blind testing has remained central in KBS validation from the earliest systems on (see, for instance, the validation of the MYCIN system (Buchanan and Shortliffe, 1984) and blind testing in medical KBS (Chandrasekaran, 1983)).

In addition to testing, KBS developers realized that rule bases could be analyzed for anomalies which are indicative of errors in the construction of the KBS, and which can lead to faulty behaviour at run-time. Commonly considered anomalies included inconsistency, redundancy, subsumption, circularity, unreachable goals, and unfireable rules. Tools to detect such anomalies were called verifiers due to the logical nature of the tests. Early verifiers, performed pair-wise comparison of rules (Suwa, Scott and Shortliffe, 1982); more sophisticated techniques including the effect of rule chaining were used by Ginsberg (1988). The use of such verifiers has been widely acknowledged as being complementary to the necessity of testing. Nowadays, validators perform a combination of verification and testing methods in order to obtain maximum evidence as to the correctness of KBS.

It is well known that software validation cannot be delayed until implementation. Otherwise, there is too high a risk that errors will be found late which may be very expensive to correct. This principle, coming from software engineering, also applies to knowledge engineering. However, most of the validation approaches developed for KBS assume to work on an implemented system (in the context of prototyping). Several authors have made proposals to include validation during the early stages of KBS development, but this does not appear to have become common practice, and there is consequently little published evidence of the practical usability of existing techniques early in development.

2.1 Dominant V&V Techniques for KBS

Currently, the dominant techniques for V&V activities can be clustered in four main groups:

- Inspection
- Static verification
- Empirical testing
- Empirical evaluation

Inspection techniques aim at detecting semantically incorrect knowledge in the KB. Inspection is performed manually, by a human who has expertise in the application domain. During development this is usually the same expert who provided the knowledge for the KB, but at some point the KB should be inspected by an expert independent of those involved in the KBS development. (Typically, this technique can be used only infrequently due to the lack of availability of experts.) Inspection is most able to identify errors in isolated KB elements: when errors come from the interaction of several KB elements — for instance, chaining of several rules — human inspectors are usually unable to detect it "by eye".

Static verification checks the KB for anomalies. An anomaly is a static pattern in the KB structure which suggest the presence of an error in the encoded knowledge. Typically, the anomaly pattern is a counterexample of a general property which should hold in the KBS (for instance, consistency). Detected anomalies need to be analyzed to determine whether they

represent an real error or just a minor defect coming from the encoding process in the selected knowledge representation. Only the most limited verification checks can be performed manually: generally this process requires computational support by automated tools. Depending on the capabilities of the verification, the checks it may perform range from a limited to an exhaustive search for anomalies in the KB. Although the properties to be checked are to a great extent domain-independent, verification tools depend on the specific semantics of the knowledge representation language used. For this reason, verifiers cannot be reused among KBS using different knowledge representation languages.

Empirical testing aims at checking KBS correctness by executing the system on sample data sets. To guarantee complete correctness, testing has to be exhaustive; that is, every potential input should be tested. This is obviously not feasible for real applications, so testing only analyzes a finite set of test data, the *test set*. The selection of the test set is crucial to the effectiveness of the testing process. Among several testing strategies, two seem to be the most effective when applied to KBS: structural and functional testing (Zualkernan, Tsai and Kirani, 1992). Structural testing aims at executing as many different KB elements as possible. Functional testing aims at checking the different KBS functions according the requirements, without considering internal structure. Real test cases are usually scarce, so test cases have to be synthesized automatically by test case generators. A final difficulty in KBS testing occurs when the application domain is so ill-defined that "correct" behaviour is not well-defined (there is no "gold standard"). In such case, it is necessary to make some definition as to what is to be considered a "correct" or "acceptable" solution for each test case; usually, the correct solution is approximated by a consensus among the opinions of several human experts.

Evaluation addresses the relation between the operational KBS and the final user. Typical evaluation issues are technical performance, acceptability, inclusion in the organization, responsibility issues, and so on. Empirical evaluation is performed by using the operational KBS either in a controlled environment (laboratory evaluation) or in the working environment (field evaluation). KBS evaluation is a human activity which is highly application-dependent.

Of these four groups, inspection and empirical evaluation methods are clearly application-dependent and they are not candidates for potential reuse across different KBS. On the other hand, verification and testing *methods* can be reused to a great extent on different KBS, even though the computational tools supporting the methods are usually bound to specific knowledge representation languages. However, even in the reusable methods, the role of human experts remain significant, because they are needed to evaluate verification and testing outputs.

2.2 Limitations of Current Approaches

While success has been achieved using the above techniques, there are still causes for concern. The presence of requirements which are hard to formalize or express without ambiguity induces different weaknesses in the V&V process. This kind of requirement is difficult to validate, and so one never knows to what extent such requirements are fulfilled. This problem becomes more difficult when different requirements coming from different users are to be integrated. Here there is an extra issue: to guarantee the internal consistency of the set of requirements. In any case, there is no rigorous way to verify the correspondence between a set of requirements and its final implementation.

The absence of formal specifications limits the capacity of KBS verification, which remains constrained to assuring that some general logical properties (such as consistency, no redundancy, no cycles, no useless objects, etc) hold in the system. These properties appear as prerequisites of a correct function of a KBS, and they should be tested. However, although they are necessary they are not sufficient, they say little about the correctness of the system behaviour. Consider a (somewhat exaggerated) example comparing knowledge engineering with numerical

programming: checking a rule base for consistency is analogous to checking in a numerical program that no computation arrives to overflow or underflow. Obviously, it is something useful to know but it has nothing to do with the correctness of the computation: one can have a consistent rule base (or a neither-overflowing-nor-underflowing numerical program) that performs totally incorrect computations, which nevertheless are consistent!

Results from testing indicate that the system performs adequately on a set of cases, and it is reasonable to expect that the system will behave in the same way on similar cases. If the test set is representative of the set of possible inputs (which in complex KBS hardly occurs), this may represent enough evidence for non-critical tasks. However, when a new case appears, we cannot be sure that the system will behave properly on it (except for trivial cases, or very specific applications where advanced testing techniques can be applied).

In summary, the main weakness of current V&V approaches on implemented KBS lies in the fact that they do not really bring confidence in the quality of the system as a whole and of the system parts. V&V methods should provide adequate answers to issues such as correctness, completeness, robustness, precision, safety, and so forth. Currently, many of these questions are answered only partially or by indirect evidence only. To overcome these defects V&V methods have to employ more precise information about the *task* the KBS is required to perform. Formal specifications can play a fundamental role in accomplishing this goal.

3 FORMAL SPECIFICATIONS FOR V&V

Strictly speaking, no verification or validation is possible without specifications, by the definition of verification (Section 1). While some of the existing verification techniques may *appear* to operate in the absence of any specifications, in reality they make reference to *implicit* specifications; for example, the techniques for verifying the consistency of knowledge bases make reference to some model of consistency, which can be considered a domain-independent requirement (that the KBS be free from inconsistency) (Preece, Shinghal and Batarek, 1992). Similarly, early validation efforts comparing a KBS against human performers made implicit reference to a requirement that the KBS should emulate such human performers (Buchanan and Shortliffe, 1984).

Once the necessity for specifications is clear, it must be said that the more precise and detailed are the specifications, the more thorough will be the V&V process. At the least-desirable extreme we have the case where the sole requirement is implicit emulation of human performers. If we have informal or semi-formal statements-of-requirements for the system, then we can devise tests to determine if the system complies with these requirements; however, because the requirements are not stated formally:

- there is no way to be sure if we have tested them completely (or to what extent we have tested them);
- they may be ambiguous, incomplete and incorrect.

The most desirable case, then, is to have requirements that are specified formally (with a well-defined syntax and semantics). Such specifications are produced by formal methods and, thus, such methods clearly have a role in the V&V process for KBS.

3.1 Formal Specifications for KBS

A growing number of formal specification languages are available — and have already been used — to specify KBS. We can categorise them roughly according to the intent of their developers:

- *General purpose* specification languages, developed in the context of conventional software engineering, for the specification of any type of software system. For example, Z (Plant and Gold, 1990) and VDM (Haugh, 1988) have been used to specify KBS.

- *Special purpose* specification languages, predominately European in origin, developed for the purpose of specifying complex knowledge-based reasoning systems. Among the best known are (ML)², DESIRE and KARL, surveyed in (Fensel and van Harmelen, 1994).

For our purposes, the second category is more attractive, not merely because the facilities of the languages are well-suited to specifying KBS, but also because these languages are better-suited to the practical needs of KBS developers. Most importantly, most of the special-purpose KBS specification languages do not expect that a fully-formal specification will be developed *a priori*; rather, they expect that the specification will be developed gradually, as a refinement of informal and semi-formal descriptions of the system: see Figure 1. Here, the informal description — typically called the "conceptual model" in various methodologies — is iteratively refined to create a more precise formal specification, with *both* descriptions undergoing gradual modification during the process. (The informal description is kept because it may be easier to communicate with domain experts using this description.) The formal description will form the basis for the implementation: later we discuss the various ways in which this can be done.

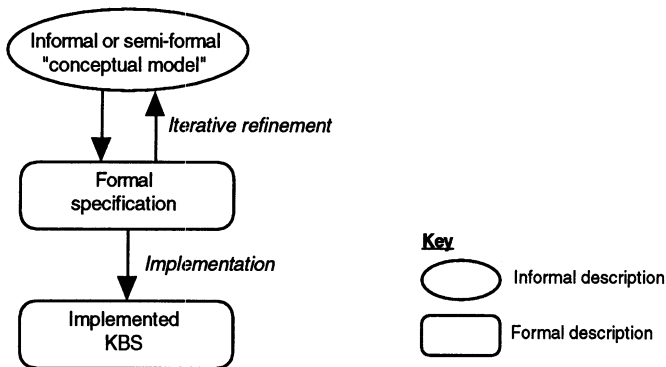


Figure 1 Using a formal specification in KBS development.

In fact, the possibilities of using formal specifications are richer than shown in Figure 1. We may want to have several formal specifications, providing different formal descriptions of the system, for example:

- At different levels of detail; for example, a specification of the system as a "black box", showing the input-output relations only, and a complementary description of the system as a "glass box", showing aspects of its required behaviour. These two views have been referred to as the *problem specification* and the *solution specification* — or as the "contract" and "blueprint" — respectively (Batarek, Preece, Bennet and Grogono, 1991).
- Of different "models" of the system; for example, the *cooperation model* may describe how the system interacts with its users, while a separate *task model* may describe the tasks the system performs on its own (Wielinga, Schreiber and Breuker, 1992).

3.2 Formal Specifications for V&V

From an idealistic point-of-view, the goal in performing verification and validation is to deliver a KBS that is as reliable as possible; more pragmatically, the goal is to deliver a KBS that is as

reliable as *necessary* — given the users' needs (Miller, 1990). The most important decision is that of *what* to verify and validate.

When formal specifications are available, they provide more opportunities for doing V&V than are available otherwise. Figure 2(a) shows what V&V can be done when the only available descriptions of the system are an informal statement of requirements, and the implementation itself.

The implementation can be verified for internal consistency and apparent completeness, and validated for (approximate) compliance with the informal requirements, using the methods described in Section 2, but that is all. The weaknesses of V&V conducted to this extent were highlighted at the end of the previous section. When formal specifications are available, however, many more V&V opportunities exist, as shown in Figure 2(b). Here, there are two typical levels of formal specification — "black box" and "glass box" — in addition to the implementation.

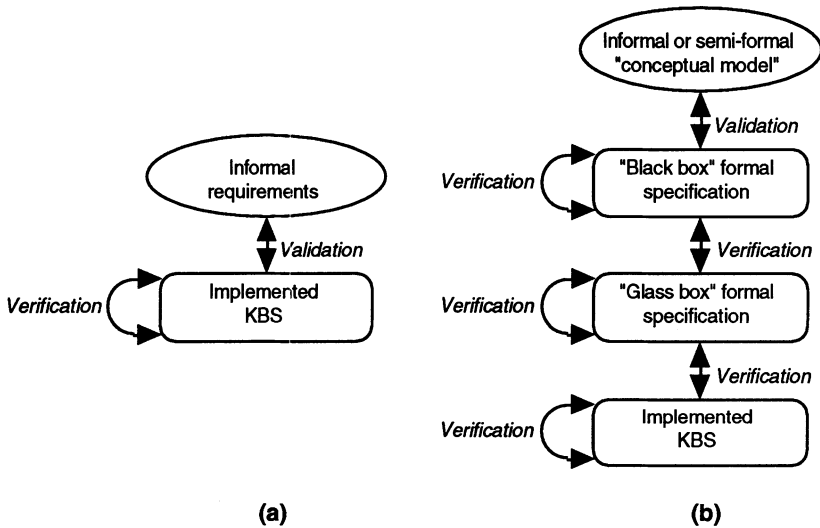


Figure 2: V&V opportunities: (a) without any formal specification; (b) with two levels of formal specification.

The availability of the multiple formal descriptions provides support for additional:

- *Intra-model V&V* (V&V of a single description of the KBS in and of itself), because the formal descriptions of the KBS can be verified and validated in and of themselves, for properties such as consistency and completeness (appropriately defined); these descriptions can also be validated manually for correctness, because the specifications are unambiguous.
- *Inter-model V&V* (V&V between different descriptions of the system), because different descriptions of the system can be verified with respect to one another; this supports validation because one level of description can be the user's requirements, made explicit.

Since the formal specifications are descriptions of the system, V&V techniques may be used to ensure the validity of the specifications themselves in the above contexts.

As described in Section 1, a formal specification provides (i) a description of the intended I/O behaviour, and (ii) a description of how this behaviour can be constructed. Some of the potential uses of a formal specification in V&V are the following:

- To give *structure* to the domain knowledge, providing a vocabulary of knowledge elements which can be used in V&V.
- To allow verification of the intended I/O behaviour in a "black box" manner.
- To allow verification of the way a solution is constructed by the system, in a "glass box" manner. The structure of the domain knowledge plus the specification of how the solution should be constructed must be sufficient to devise V&V techniques specific to the task performed by the system, to provide the necessary level of confidence in the reliability of the system.
- Some languages, such as KARL (Fensel and van Harmelen, 1994) produce executable specifications. They can be used to test the system against the ideal behaviour provided by the specifications.

The various V&V techniques described in Section 2 can — and in some cases must — be redefined in this context. The question of verifying the specification itself is still open, as few of the existing special-purpose KBS specification languages have well-developed proof techniques at present. If general-purpose specification languages are used, then their existing proof techniques apply, but these have been found to entail a great deal of labour when KBS are specified (Haugh, 1988). However, extended versions of the existing V&V techniques for KBS become possible: we consider each of the four techniques described in Section 2.

Inspection

In intra-model V&V, at the specification level, it is often argued that the unambiguous nature of formal specifications permits easier direct validation via inspection by customers and — in the case of KBS — domain experts. This is made difficult in practice when unfamiliar and unintuitive notations are employed for the formal descriptions, although this difficulty may be alleviated to some extent because the declarative nature of much of the specification (for example, logical rules) is more comprehensible to non-programmers than procedural descriptions. We believe, however, that it will usually be more practical to employ the semi-formal conceptual model in this capacity as a more intelligible (though not unambiguous in and of itself) version of the specification.

The conceptual model provides an implementation-independent description of the intended task, and therefore can be inspected by domain experts without the burden of implementation details. At this level, inspection aims at evaluating the knowledge quality and completeness with respect to the intended task. If necessary — due to ill-definition or ambiguity — formal specifications can complement or substitute the conceptual model. In this case, domain experts would likely require assistance from the knowledge engineers to deal with formal languages.

In inter-model V&V, the high degree of structure provided by a formal specification can be used manually to check for the detection of corresponding structures in the implemented system. The conceptual model may play a role here also: at the implementation stage, one can check the correspondence between the conceptual model and the final implementation of the system.

Static Verification

In intra-model V&V at the specification level, if a formal specification is to be verified according to the properties for static verification (consistency, redundancy, etc), then these properties need to be defined in terms of the formal specification language. Once this is done, however, the reusability of this approach is high, because the same anomalies can be checked for in any specification written using the language. Furthermore, formal specification languages provide additional opportunities for verification by some of the techniques, as a benefit of their features. For example, the modular architecture and declaration of hierarchies of types (sorts) provided by

languages like (ML)², permit additional properties to be checked in static verification (for example, violation of modularity, and type mis-matches).

Static verification can also be employed to detect structural nonconformances between specification levels and implementation. This provides an automatic (and hence more reliable) version of the inter-model inspectional verification discussed above. Recall that, in addition to the I/O behaviour, a formal specification describes how this behaviour can be achieved. This implies that a formal specification provides, to some extent, elements of the KBS structure. After implementation, these elements can be checked to verify whether their functionality meets their specification. This is an important step towards constructing correct and reliable KBS because (i) verification is no longer limited to general purpose properties; it can check domain-dependent properties, and (ii) verification can be made of specific parts of the KBS structure, providing evidence of correctness for these parts.

Empirical Testing

In an idealised view of the development process with formal specifications, testing is typically considered to be unnecessary, because the specifications are intended to be self-evidently correct, and all subsequent descriptions are shown to be compliant with the previous specifications. Although the presence of formal specifications may suggest that testing is unnecessary, this view is currently too extreme and impractical, and is likely to remain so for the foreseeable future. The translation from specifications to implementation is performed manually, and some errors can be introduced. Therefore, testing is still necessary, but the existence of formal specifications can help greatly in the selection of the test set. A formal specification provides the intended I/O behaviour, so it contains all the information needed to perform functional testing. On the other hand, the level of structure in formal specifications can be used to support structural testing.

Furthermore, with the existence of formal specifications, it becomes possible to test the specifications directly; such testing can take the conventional form of running test cases on an executable specification, or can take the form of proving properties (such properties corresponding to "test cases") of the specification. In order to do this, proof techniques must have been defined for the specification language.

Empirical Evaluation

This aspect of V&V would seem at first to benefit little from the use of formal specifications. However, a number of possibilities exist. For example, it is possible to use a formalised model of the required cooperation between KBS and users to check for potential system integration problems in a more principled way — and at an earlier stage in development — than would otherwise be possible. Secondly, subjective requirements such as "user friendliness" can be approximately expressed by means of pseudo-formal specifications, and then become amenable to V&V methods other than experimental evaluation (Laurent, 1992).

3.3 Outstanding Issues

While formal specification techniques offer clear benefits to V&V, there are a number of unclear issues at present.

Nature of the Development Process

One question concerns the nature of the development process incorporating formal specification with V&V. Should a transformation approach be adopted, wherein the specification is gradually refined into an implementation, as shown in Figure 3(a)? This is attractive from the V&V point-of-view because it is easier to promote and control validity throughout the process: all verification is performed directly upon the formal specification (benefitting from its well-defined syntax and semantics), and correctness is then assured by the transformation process.

However, it is harder to ensure that the implementation will be efficient when the transformational approach is taken, and the progression through successive levels of refined specification may place a heavy burden on developers. This contrasts with the more conventional approach shown previously in Figure 2(b), where an implementation is crafted with the intention that it will comply with the specification, without actually being derived from it. V&V techniques are used to establish the compliance between the descriptions.

Completeness of Formal Specifications

Another important question concerns the completeness of the specifications; that is, whether the specification should (and, in complex cases, whether it *can*) describe every aspect of the system. One approach to development is to aim towards building a formal specification which is a complete description of the KBS. This can then be transformed (as in Figure 3(a)) or constructed into an implementation.

An alternative approach is to specify different aspects of the system as independent "mini-models", against which the implemented system can be verified (Bellman, 1990) — see Figure 3(b). There may be a strong practical reason for choosing the latter approach: when the KBS is complex and ill-structured, requiring a great deal of knowledge acquisition, analysis and refinement (typically supported by exploratory prototyping) before a reasonably complete version exists. However, transformational implementation is no longer possible because the specification is not complete.

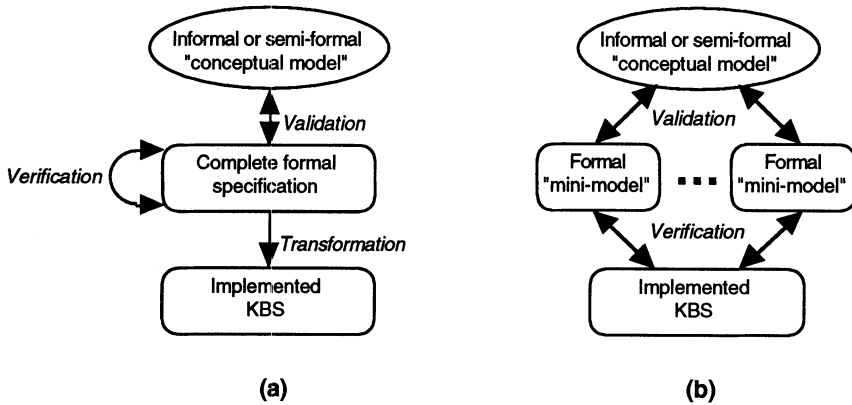


Figure 3: Using formal specifications for V&V: (a) transformation from complete formal specification; (b) use of "mini-models".

4 CONCLUSIONS

In knowledge engineering, the V&V activity is a kind of "watching eye" aimed at detecting deviations between what is intended to be built and the artifact being built. To perform this task, V&V requires information about the intended KBS. The more information that is available, and the more precise that information, the better the validation process will be. This will have a direct impact in the quality of the final KBS.

Conceptual models of the intended task appear to be essential in KBS construction. These models play a fundamental role in the validation process, because they act as the reference to compare against at any stage of KBS development. Conceptual models allow one to structure the domain knowledge, identifying domain-dependent properties which can be effectively tested. These properties can provide direct evidence of the different dimensions to be evaluated in V&V. Importantly, this checking can be made at any stage in the development process, so the difficulties to include V&V activities during KBS construction are removed by this approach.

Conceptual models can support the development of formal specifications. The presence of formal specifications eliminates one of the main weaknesses in V&V process: the use of vague or ambiguous requirements. Formal specifications can be used to verify the system, and to define precisely its boundaries (one of the classical problems in KBS validation). In addition to describing the I/O behaviour, formal specifications define how this behaviour can be achieved. This brings more information for V&V, because a formal specification provides, to some extent, elements of the KBS structure. These elements can be effectively used in the V&V process, which is no longer forced to be a "black box" testing process.

The construction of conceptual models and their translation into formal specifications is neither an easy nor currently a complete task. Some typical KBS tasks remain elusive to formal specification, and the specification construction can be a costly process. Topics for immediately-necessary future work include:

- making concrete the proposals for using formal specifications with the various V&V techniques, and evaluating the effectiveness of doing so;
- investigating the utility and practicality of different approaches to formal specification (e.g. "mini-models" versus complete specification);
- development of development processes in which all of the techniques can be applied and managed effectively;

These issues remain open in knowledge engineering, but they should not be obstacles to the development and use formal specifications for KBS construction. Their benefits will improve the quality of KBS in the near future.

5 REFERENCES

- Batarekh, A., Preece, A.D., Bennett, A. and Grogono, P. (1991) Specifying an expert system. *Expert Systems with Applications*, 2, 285-303.
- Bellman, K.. (1990) The modeling issues inherent in testing and evaluating knowledge-based systems. *Expert Systems with Applications*, 1, 199-215.
- Buchanan, B. and Shortliffe, E. (1984) The problem of evaluation, in *Rule-Based Expert Systems: the MYCIN Experiments of the Stanford Heuristic Programming Project*, 571-588, Addison-Wesley, Reading MA.
- Chandrasekaran, B. (1983) On evaluating AI systems for medical diagnosis. *AI Magazine*, 4, 34-37.
- Fensel, D. and van Harmelen, F. (1994) A comparison of languages which operationalise and formalise KADS models of expertise. *Knowledge Engineering Review*, 9, 105-146.
- Ginsberg, A. (1988) Knowledge-base reduction: A new approach to checking knowledge bases for inconsistency & redundancy, in *Proc. 7th National Conference on Artificial Intelligence (AAAI 88)*, vol. 2, 585-589.
- Haugh, J. (1988) The application of formal specification techniques to knowledge-based system development, in *UK IT 88 Conference Publication*, 95-98, Information Engineering Directorate, London.

- Laurent, J-P. (1992) Proposals for a valid terminology in KBS validation, in *Proc. 10th European Conference on Artificial Intelligence (ECAI 92)*, 829-834, John Wiley & Sons, Chichester.
- Miller, L.A. (1990) Dynamic testing of knowledge bases using the heuristic testing approach. *Expert Systems with Applications*, 1(3):249-269, 1990.
- O'Keefe, R.M., Balci, O., and Smith E.P. (1987) Validating expert system performance. *IEEE Expert*, 2, 81-90.
- Plant, R.T. and Gold, D. (1990) Increasing expert system reliability through the use of a formal specification, in *AAAI-90 Workshop on Knowledge Based Systems Verification, Validation and Testing*, AAAI, Menlo Park, CA.
- Preece, A.D. (1990) Towards a methodology for evaluating expert systems. *Expert Systems*, 7, 215-223.
- Preece, A., Shinghal, R., and Batarekh, A. (1992) Principles and practice in verifying rule-based systems. *Knowledge Engineering Review*, 7, 115-141.
- Suwa, M., Scott, A. and Shortliffe, E. (1982) An approach to verifying completeness and consistency in a rule-based expert system. *AI Magazine*, 3, 16-21.
- Turing, A. (1950) Computing machinery and intelligence. *Mind*, 59, 236-248.
- Wielinga, B., Schreiber, A. and Breuker, J. (1992) KADS: a modelling approach to knowledge engineering. *Knowledge Acquisition*, 4, 5-54.
- Zualkernan, I., Tsai, W-T. and Kirani, S. (1992) Testing expert systems using conventional techniques, in *Proc. 16th Annual Computer Software and Applications Conference*, 320-325.