Ontologies in Legal Information Systems;
The Need for Explicit Specifications of Domain Conceptualisations

Trevor J.M. Bench-Capon and *Pepijn R.S. Visser*

LIAL - Legal Informatics at Liverpool
Department of Computer Science
University of Liverpool
PO Box 147 Liverpool L69 7ZF
United Kingdom
{tbc, pepijn}@csc.liv.ac.uk

Abstract

In this paper we discuss the role of explicit specifications of domain conceptualisations - now popularly called "ontologies" - in legal information systems. We describe the advantages that accrue from producing an ontology for such systems, as well as the ontologies so far developed which are directed at the legal domain. We then illustrate how taking an ontological perspective can give insight into what is common and what is disparate in apparently different approaches. We conclude by offering some findings as to the nature of ontologies for legal information systems, and some ideas concerning the creation of a library of legal ontologies.

1. Introduction

In the last five years there has been a growing recognition of the need to create explicit specifications of how knowledge in a domain is conceptualised. Such an explicit specification is now normally termed an "ontology" in the literature, following Gruber (1992). Whether this term is helpful or not - certainly it cannot be understood in the traditional sense with which it has been used for centuries in philosophy (*e.g.*, Kant, 1781) - is not something we will discuss here (*e.g.*, Guarino and Giaretta, 1995): what is important is that the term is widely used in computer science. We shall therefore use the term "ontology" in Gruber's sense throughout this paper.

ICAIL-97, Melbourne, Australia © 1997 ACM 0-89791-924-6/97/06.$3.50

In section 2 we will describe the various motivations for producing ontologies. Section 3 will provide an illustrative example. Section 4 will look at two ontologies that have been produced for legal information systems. Section 5 will look at some previous work from an ontological standpoint, and show how this can give insight into existing systems. We conclude this paper in section 6 by presenting our findings and suggestions concerning the creation of a library of legal ontologies.

2. Motivation for Using Ontologies

Since Gruber is the originator of the term in the field of AI, we will first look at his motivation. In doing so we will make more precise the sense in which he uses the term.

(1) *Knowledge Sharing.* Gruber's motivation for using ontologies is knowledge sharing. His work has been done in connection with the ARPA knowledge sharing initiative (Gruber, 1993), which has as its vision the intercommunication of a federation of knowledge bases, which are able to pool their knowledge to solve problems. In such a context a system-independent understanding of the knowledge contained in these disparate systems is crucial. Gruber's approach is express this common understanding as based on "ontological commitments", which he says are "agreements to use the shared vocabulary in a coherent and consistent manner" (Gruber, 1995). A conceptualisation underlies any formally represented body of knowledge and hence any knowledge system is committed to a conceptualisation, although this is too often left implicit.

By conceptualisation we understand the objects, concepts and other entities that are assumed to exist in some area of interest and the relationships which hold between them (Genesereth and Nilsson, 1988). If this conceptualisation remains implicit, then the knowledge cannot be shared with any confidence, since the assumptions may not match. The idea of an ontology, in the context of knowledge sharing, is thus to make explicit what must be committed to if a common understanding of represented knowledge is to be achieved. With respect to knowledge sharing, the role is similar to that of a schema in databases. Since the database is meant to hold data to be used by a variety of programs the authors of those programs must be in a position to know what data is available, and what form the data is held in (*e.g.,* type information, range constraints). Just as the schema is essential for the sharing of data between application programs, the ontology is essential for sharing knowledge between different knowledge systems.

(2) *Verification of a knowledge base.* Traditionally the acceptability of a knowledge has been ascertained by testing - if the behaviour is correct the knowledge base is assumed to be correct. This corresponds with the validation of traditional systems. But traditional systems also typically see verification as of importance, whereby the program is compared with its specification to ensure that it has been built correctly. This, of course, requires a specification to verify against, and this has usually been lacking in knowledge systems. As a replacement it has been argued that examination of declarative code can provide verification enough, but this is so only if the implicit ontological commitments of the code are shared by the reader. Without an explicit specification of the conceptualisation, there can be no assurance that this is so.

(3) *Software engineering considerations.* Any program requires proper documentation, to guide both end-users and any future maintainers of the program. Standards for documenting knowledge systems have been slow to emerge. This has often led to worries about assumptions implicit in the system. For example, the discussions about the interpretation of is-a links (*e.g.,* Brachman, 1986)) in which it can be unclear as to whether they are intended to be exhaustive, exclusive, and whether they are used consistently. An ontology provides much of the documentation which is required and provides definitive answers to questions of this sort.

(4) *Knowledge acquisition.* Too often knowledge acquisition is a somewhat haphazard process, guided by the experts and the skill of the knowledge engineer in eliciting knowledge from them. If, however, the conceptualisation is explicit, the knowledge engineer has a framework with which to guide the knowledge acquisition: the kind of knowledge that must be acquired is described from the outset rather than expected to emerge from discussion. Of course, this means that there must be a prior stage in which the domain is conceptualised (often with the assistance of an expert): such a stage is explicitly incorporated in certain knowledge acquisition methodologies such as MEKAS (Bench-Capon *et al.*, 1993).

(5) *Knowledge Reuse.* Given a conceptualisation of a domain we can - if it is explicit - exploit it in the design of systems in the same or related domain. In the case of the same domain it can simply be adopted and used to supply the vocabulary to design another system, or it may be necessary to refine the conceptualisation to take it to a greater level of detail which adapts and extends it for a new application in the domain (Visser and Bench-Capon, 1996). In a related domain we may wish to use only the higher levels of the conceptualisation and refine them in a different direction. In any event the "ontology" of one application will be a reusable component on which to ground the design of the new application.

(6) *Domain-theory development.* In some domains there may be unanimity on what is an appropriate conceptualisation. In others - and law is certainly such a domain - there may be radical differences in the way in which the domain is conceptualised. Later in this paper we will briefly discuss two ontologies for law which appear to have little point of contact. If we are to analyse, compare, and perhaps even reconcile, different conceptualisations, we must have an acceptable starting point, and this can only be done with an explicit and unambiguous specification of the conceptualisation. A current handicap to theory development in AI and Law is the difficulty of making comparisons between the work of different groups. Much work in AI and Law embodies certain fundamental assumptions about the nature of legal knowledge; for example we may see law as conceived of as definitions (*e.g.*, Sergot *et al.*, 1986), or as cases (*e.g.*, Ashley, 1990). Often these assumptions remain rather more implicit than the works cited above, and then it becomes hard to discover just what is intended by the work. Certainly the difference goes deeper than the representation formalism used: it is possible to represent cases using rules and definitions using frames. An ontology would greatly facilitate fruitful discussion and comparison of different approaches.

All, indeed any, of the above would provide a powerful reason to provide an ontology when producing a knowledge system. All of them apply with equal or greater force when the knowledge system is in the domain of law: the inter-relation of law makes it a natural area for knowledge sharing; the importance of legal decisions argue for a high level of verification; the rate of change of law argues for readily maintainable systems, a well-known software engineering problem; knowledge acquisition is no less a problem in law than in other domains; the similarity of different branches of law urges the design of re-usable frameworks, and the lack of fundamental theoretical agreement suggests that we should reap whatever insights in the way of domain theory development are available. Moreover, the lesson of work in AI and Law is that there is far too little interaction between different workers in the field, and far too little building of the work of others - a situation which will only be remedied if researchers make their fundamental conceptualisations explicit. In the next section, before discussing work done in the AI and law area, we will discuss a simple example to make some of the above points more concrete.

3. An Example

UK Social Security Law has a concept of "pensionable age", attained at 65 in the case of a man and at 60 in the case of a woman. In a legal knowledge system, which uses logic programming as its method, this might be represented as:

```
C1  pensionable_age(X) :-
        sex(X, male), age(X, A), A >= 65.
C2  pensionable_age(X) :-
        age(X, A), A >= 60.
```

Alternatively, we might represent this as:

```
C3  pensionable_age(X) :-
        sex(X, male), age(X, A), A > 64.
C4  pensionable_age(X) :-
        sex(X, female), age(X, A), A > 59.
```

If we make the right assumptions these two representations of the concept are equivalent. But, of course, if these assumptions are violated, differences emerge.

1. C1 and C2 supposes that anyone who fails *sex(X, male)* will be female. This embodies two assumptions: that the second argument of sex can only be either male or female, and that *sex(X, male)* will fail only if *X* is female, not, for example, because, the sex of *X* is unknown.
2. C3 and C4 rely on the assumption that *A* will be bound to an integer. If ages were recorded as reals to express part years of completed life, the comparisons in C3 and C4 would be incorrect.
3. C1 and C2 rely on a particular execution strategy, which means that C2 will be reached only after C1 has failed. Thus C2 contains the implicit condition that not *sex(X, male)*.

All of the assumptions described in 1 - 3 are eminently plausible, but need not be observed. An ontology would, however, make it clear, that in the case of C1 and C2 the second argument of sex is an enumerated type with two values and that in the case of C3 and C4 the second argument of age is an integer. Without such information showing concurrence of ontological commitments it would be folly to use one of the formulations in the context appropriate to the other. We note that the third point may not be addressed by an ontology of the legal domain unless it insists in some way that implicit negations are made explicit.

There is another point here, however. The form of representation used implies a commitment to a certain conceptualisation at a high level. In the particular example, a concept, pensionable age is defined, by giving a pair of sufficient conditions, which are intended to be taken together as supplying a necessary condition. The system is intended to be executed by deducing the applicability of the defined predicates from the applicability of certain more primitive predicates. Whilst the lower level specification - at the level of typing - is most crucial for the knowledge sharing and verification and validation benefits to accrue, this higher level specification is important for the advantages listed as 3 - 6 in section 2. We need therefore to be aware of the different levels of abstraction at which ontologies may be specified, and to recognise their utility at various of these levels.

4. Ontologies in AI and Law

So far the field on AI and Law has produced two substantial ontologies: the functional ontology of Valente (1995) and the frame-based ontology of Van Kralingen (1995) and Visser (1995). Both are presented in a reasonable amount of detail and both have been formalised in the same language for describing ontologies, ONTOLINGUA (Gruber, 1992).

4.1 The Functional Ontology of Valente

Valente's ontology of law is based on a functional perspective of the legal system. The legal system is considered as an instrument to change or influence society in specific directions by reacting to social behaviour. This main function can be decomposed into six primitive functions, each corresponding to a category of primitive legal knowledge:

(a) *normative knowledge*; which describes states of affairs which have a normative status, such as forbidden or obligatory. Note that it is these states of affairs which are considered to be the objects of the deontic modalities, and the actions which realise them derive their normative status from them.
(b) *world knowledge*; which describes the world that is being regulated, in terms that are used in the normative knowledge, and so can be considered as an interface between (non-legal) common sense knowledge and normative knowledge.
(c) *responsibility knowledge*; this is the knowledge which enables responsibility for the violation of norms to ascribed to particular agents
(d) *reactive knowledge*; which describes the sanctions that can be taken against those who are responsible for the violation of norms
(e) *meta-legal knowledge*; which describes how other legal knowledge should be reasoned with. For example it would include principles such as *lex specialis* to assist in resolving conflicts in legal knowledge.
(f) *creative knowledge*; which states how items of legal knowledge can come into being and cease to be.

This ontology forms the basis of a system ON-LINE (Valente, 1995) which Valente describes as a "Legal Information Server", the chief feature of which is the storage of legal knowledge as both text and as an executable analysis system, interconnected through a common expression within the terms of the functional ontology. The key thrust of this conceptualisation is thus to act as a principle for organising and relating legal knowledge, particularly with a view to conceptual retrieval.

Two limitations are noted by Valente. The first is practical - that performing the modelling required to follow through this conceptualisation is very resource intensive. Although the Ontolingua description of the different kinds of legal knowledge seems relatively complete, the domain model constructed within this framework for the ON-LINE system, is rather restricted. The second is theoretical - to what extent does the ontology generalise to different varieties of law. Valente writes:

> "While it is expected that the ontology is able to represent adequately legal knowledge in several types of legislation and legal systems, this issue was not yet tested in practice" (p.175).

An additional question relates to whether the functional ontology corresponds to anything that might be found in legal theory. This is a fair question since Valente himself argues that ontologies can and should serve as the "missing link" between legal theory and AI and Law-working as "a neutral and problem independent medium whereby legal theoretical ideas may be expressed" (p.45). Valente writes with respect to its cognitive validity:

> "The model of legal reasoning embedded in this ontology is very unlikely to be cognitively valid and is probably counter-intuitive in several respects to both the average reader and the legal theorist. ... the task of the ontology is neither to match these intuitions, not to explain empirical evidence as to how judges reason, but to find the most parsimonious structure for representing and reasoning with legal knowledge in a valid way" (p.77).

This last quotation suggests that the purpose of the ontology is largely computational and that it is by its value as the basis for a computer system that it should be judged, rather than by its correspondence to legal theory.

4.2 The Frame Based Ontology of Van Kralingen and Visser

The second ontology we will discuss was motivated by a desire to improve development techniques for legal knowledge systems, and in particular to enhance the reusability of knowledge specifications by reducing their task dependency. This ontology distinguishes between an ontology which is intended to be generic to all law, and a statute-specific ontology which contains the concepts relevant to a particular legal domain. The two are related: the statute-specific ontology can be seen as a refinement of the generic ontology which fits it for a particular domain.

The generic ontology divides legal knowledge over three distinct entities: norms, acts and concept descriptions. For each of these entities the ontology defines a frame structure that lists all attributes relevant for the entity. Both the division of legal knowledge into these three entity types, and the attributes identified for these entities are held to be true of any area of law. However, modelling a legal sub domain also involves resolving numerous ontological matters, which are concerned with the particular content of legal questions in the specific domain under consideration. For instance, is it necessary to distinguish between male and female employees in the Dutch Unemployment Benefits Act? Such questions motivate the distinction between the legal and the statute-specific ontology. Essentially, the statute-specific ontology consists of predicate relations that are used to instantiate the frames for norms, acts and concept descriptions, in a manner appropriate to the specific domain. Van Kralingen and Visser argue that the generic component can be re-used in any legal domain, whereas the statute-specific ontology can be re-used only for different tasks within the same domain.

This ontology has been used as the basis for the system FRAMER which addresses two applications in Dutch Unemployment Benefit Law (Visser, 1995), one involving a classification task determining entitlement to Unemployment Benefit and the other a planning task, determining whether there is a series of actions which can be performed to bring about a certain legal consequence.

Similar questions can be asked about this ontology as were asked about that of Valente. First, it is again true that building the ontology is resource intensive: however one can say that this may be compensated for in the generic nature of the upper levels of the ontology. Moreover, building a proper computer system just *is* resource intensive. The situation is not very much different from that of producing a proper entity-relation model for a database system. Second, one can speculate as to how generic the ontology really is: the domain of law to which it was applied is very definitional in nature, and it would need to be applied to other domains to prove its generic nature. Third, we can again inquire as to its relation to legal theory. It is certainly based on a legal-theoretical analysis of the building blocks of law (Van Kralingen, 1995, ch. 2) and is intended at least to be concordant with legal theory.

4.3 A Brief Comparison of the Two Ontologies

The most obvious thing to say about the two ontologies described above is that they are very different. At the highest level they diverge immediately. This is because of the different perspectives from which they begin their conceptualisation: Valente seeks a functional decomposition of the legal system considered qua system, whereas Van Kralingen and Visser seek a set of building blocks from which they can construct law qua body of knowledge. Given such divergence, we might wonder if any further comparison will be fruitful. We can, however, if we delve deeper we find some instructive similarities. In the ontology of Van Kralingen and Visser we find defined as relations "event-qualification" and "process-qualification". The intention of these relations is to relate a physical event with an institutional event. This is a crucial step in their ontology since it enables physical acts (*e.g.*, A kills B), to be classified in terms of institutional acts (*e.g.*, A murders B, or alternatively, A manslaughters B) which are what tend to be used in norm descriptions. It is precisely these relationships between the physical and the institutional descriptions that are the subject of Valente's category of "world knowledge". Thus we can see that in both conceptualisations the transition from the physical description to the institutional description is of crucial importance. In the case of normative status, both ontologies define it as a function, although for Valente it maps from situations to a normative status, whereas in Van Kralingen and Visser it maps from a(n) (applied) norm to a normative status. There remains, however, a similarity in that we can go in the latter ontology to a situation via the act prescribed in a breached norm which has a slot giving the post-conditions of the act. A difference remains, however, in that in Valente's ontology there are three flavours of normative status, allowed, disallowed and silent, whereas Van Kralingen and Visser subsume both allowed and silent under "not breached".

This paper is not the place to go into an exhaustive description of the similarities and differences between the two ontologies. For a more detailed comparison, see Visser and Bench-Capon (1997). The brief remarks above, however, do illustrate the value of having the explicit specification of the conceptualisation. From the informal description of the two ontologies one might well conclude that there was an unbridgeable chasm between the two. The formal ONTOLINGUA specifications, however, show that there are in fact considerable similarities, which could form the basis of translation rules which would permit knowledge sharing, as well as the kind of subtle differences illustrated by the variations in the types of normative status allowed.

5. Taking an Ontological View of Existing Knowledge Bases

Both the above ontologies were developed as a precursor to building a knowledge base for a particular application, in approved software engineering style. There is, however, also value in thinking about what an ontology to support an existing knowledge base would be like, so as to gain insight into and understanding of some of the classic approaches adopted in AI and Law (*cf.* Moles and Dayal, 1992, p.205). In this paper we will illustrate this by considering the logic programming approach, most famously exemplified

in the formalisation of the British Nationality Act (Sergot *et al.*, 1986).

The knowledge base here comprises a set of definite clauses of the form $A \leftarrow B_1 \ldots B_n$, which are to be read as "A is true if all of $B_1 \ldots B_n$ are true. The motivation for representing law in this way is that

> "many kinds of legislation are essentially definitional in nature, and that extended definite clauses provide a simple and natural formal language for expressing such definitions precisely" (Sergot, 1991, p.36).

The significance of "extended" here is that the right-hand sides of the clauses are to be allowed to include literals of the form "not B_i", and that these negations are to be interpreted using negation as failure. This has the consequence that while the reading of a definite clause suggests that a sufficient condition for its head is provided, the set of clauses for a given head taken together must be interpreted as supplying a necessary condition for the truth of that head. Attempts have been made (*e.g.*, Kowalski, 1989) to enable discrimination between the case where this is desirable and where it is not, but in the British Nationality Act program, the only-if parts were taken as supplied.

The above suggests that legal knowledge is conceived within this approach as being a set of definitions, and that the definition of a term is conceived of as comprising a set of extended definite clauses, held to supply individually sufficient and jointly necessary conditions for the obtaining of that term.

A full reconstruction of the ontology would require a specification of all allowable predicates (including their possible instantiations), but we will not go into that level of detail here. What we will note, however, is that the predicates form two natural groups: those which have a definition in the knowledge base, and those which do not. The latter group can be divided into those to be supplied by the system (*e.g.*, today's date), and those to be supplied by the user in response to questions (*e.g.*, date and place of birth). Whether a predicate has a definition in the knowledge base depends on whether a definition can be found in the legislation being formalised.

This view of the conceptualisation indicates several things. First, it indicates that the conceptualisation is not intended to be general, but to apply only to the "many kinds of legislation that are essentially definitional in nature". Second, we can see that it operates at the level of institutional facts (Searle, 1964): for the most suitable legislation (of which the BNA is an example) the predicates put to the user will be physical facts, in which case the program resulting from the formalisation can be seen as a useful way of deriving institutional facts form physical ones. For other legislation the undefined terms will themselves be institutional facts, defined perhaps in other legislation, or through case law, and here - unless the user is in a position to decide upon them - the utility of the program is much lessened. Where we attempt to extend the formalisation to include some representation of expertise so as to assist in the resolution of these predicates, as in, for instance, Bench-Capon (1991), we depart from the original conceptualisation, and need to recognise this departure. Thirdly, we must note the restrictive notion of a definition that is employed. Recognition that this is not always adequate (*e.g.*, to deal with definitions in terms of counterfactuals) led to other extensions (*e.g.*, Routen, 1989). The above also helps us to evaluate criticisms such as that made by McCarty (1995). He writes:

1. "Legal concepts cannot be adequately represented by definitions that state necessary and sufficient conditions. Instead legal concepts are incurably `open-textured'.
2. Legal rules are not static but dynamic ... Thus the important process in legal reasoning is not theory application, but theory construction.
3. In this process of theory construction, there is no `right answer'. However, there are plausible arguments, of varying degrees of persuasiveness, for each alternative version of the rule in each new fact situation."

In general we must agree. But since the BNA conceptualisation is intended to apply only to terms explicitly defined in legislation, it might be argued that what he has identified is a limitation of the applicability of the approach rather than a refutation of it with respect to the law to which it is intended to apply. McCarty offers an alternative conceptualisation in which there are three components to a legal concept: a set of necessary conditions (which may be empty); a set of exemplars representing sufficient conditions, and a set of transformations expressing the relationships between the set of exemplars. The second and third sets are open and extensible; this is supposed to accommodate open texture and to incorporate the dynamic aspect since "The application of a concept to a new factual situation automatically modifies the concept itself" by extending the set of exemplars, and/or the set of transformations. With this formulation, we could, however, construe McCarty not as arguing for the abandoning of the BNA style approach, but for its extension, since it supplies his first component, but not the other two.

Attempting to reconstruct ontologies in this way gives us a better understanding of the approaches and helps to discriminate amongst apparently similar approaches, such a the pure BNA approach and an approach which additionally allows clauses representing heuristics or decisions in particular cases. It also allows us to see relationships between apparently opposed approaches.

6. Conclusion

In this section we will present a number of findings and indicate directions which we feel future research could possibly take.

1. Ontologies are designed for particular purposes. Assessing the adequacy or suitability of an ontology can only be done given the purpose the ontology is created for. The criteria which an ontology must fulfil in order to provide the basis for knowledge representation are far stricter - with respect to completeness and the detail required - than one which is merely meant to characterise an approach to legal knowledge systems for the purpose of contextualising work (section 4).

2. There is no agreement on what exactly should be specified in a(n) (legal) ontology, nor on the level of detail an ontology should be specified at. Given the variety of purposes for which an ontology might be wanted (in terms of tasks, and subdomains) we should expect substantial differences in what is presented as an ontology. The important thing is to be clear as to what the ontology is for, and to avoid using one for a purpose for which it was not intended (section 2 and 3).

3. Different authors create substantially different conceptualisations of the legal domain even when their aims are rather similar. This is important, as is avoiding the temptation to ask which is the right conceptualisation. What is needed is rather for the ontologies to be sufficiently clearly stated that it is possible to relate them and to come to understand what is motivating the differences. We must recognise that there is no agreement on the basic conceptualisation(s) of the legal domain, and explore what is offered by different perspectives (section 4).

4. Ontologies provide a useful basis for comparison and analysis of different approaches in AI and Law research. A similar claim has been made by Moles and Dayal (1992 p.17) in arguing that system developers should be studying the domain-theoretical paradigms with which they design their systems (section 5).

5. There is a trade off between domain reusability (can we use the ontology for different legal tasks and (sub)domains?) and epistemological completeness (do we cover all relevant terms?). At the lower level of detail the ontology must make commitments to a particular task and the information available to those performing the task (section 4).

6. The distinction between a statute-specific and a generic legal ontology is useful and necessary.

The former is important to express the basic conceptualisation of the domain, the latter to provide the engineering benefits required by particular applications (section 4).

The research illustrates that the same domain can be conceptualised in different ways. Each way of conceptualising the domain yields different ontologies, none of which is necessarily better than the others. We believe that future research should focus on a modular approach to building legal ontologies. To achieve this, we list a number of suggestions for future research which we deem fruitful.

- Research should not concentrate on creating *one* ontology of the legal domain but on the creation of a library that contains several *dedicated ontologies*. That is, different ontologies should be made for different purposes.
- The ontology library should contain both *competing ontologies* (*viz.* ontologies designed for roughly the same purpose but with different commitments) as well as *supporting ontologies* that can be combined to create a composite ontology.
- The library should contain ontologies at different abstraction levels, such that detailed ontologies can be used in combination with more abstract ontologies. As an example, we mention the legal ontology and the statute-specific ontology as recognised by Van Kralingen and Visser.
- The library should contain ontologies for both the world assumed by the legal system (*viz.* the Legal Abstract Model as described by Valente) and the physical view in Van Kralingen and Visser and the legal qualifications of this world (*viz.* the institutional view).

Such a library would provide an invaluable resource for the AI and Law community, in particular by giving a foundation for the exchange of results in a way that would promote integration with diverse research. Further, it would provide a useful testbed on which to evaluate the claims made for ontologies in the general AI community.

References

Ashley, K,D. (1990). *Modeling Legal Argument*, MIT Press, Cambridge, Massachusetts, USA.

Bench-Capon, T.J.M. (1991). Practical Legal Expert Systems: the Relation Between a Formalisation of Law and Expert Knowledge, Bennun and Narayanan (*eds.*), *Computers, Law and AI*, Ablex, pp.191-201.

Bench-Capon, T.J.M., F.P. Coenen, H.S. Nwana, R.C. Paton, and M.J.R. Shave (1993). Two aspects of the Validation and Verification of Knowledge Based Systems, *IEEE Expert*, Vol.8, No.3, pp.76-81.

Brachman, R.J. (1986). I Lied About the Trees, *AI Magazine*, Vol.6, No.3, pp.60-93.

Breuker, J.A., and W. van de Velde (1994). *CommonKADS Library for Expertise Modelling, Reusable Problem Solving Components*, J.A. Breuker, and W. van de Velde (*eds.*), IOS Press, Amsterdam, The Netherlands.

Genesereth, M.R., and N.J. Nilsson (1988). Logical Foundations of Artificial Intelligence, Morgan Kaufmann Publishers, California, USA.

Gruber, T.R. (1992). *ONTOLINGUA: A Mechanism to Support Portable Ontologies*, technical report, Knowledge Systems Laboratory, Stanford University, California, USA.

Gruber, T.R. (1993). A Translation Approach to Portable Ontology Specifications, *Knowledge Acquisition*, Vol.5, pp.199-220.

Gruber, T.R. (1995). Toward Principles for the Design of Ontologies Used for Knowledge Sharing, *Int. Journal of Human-Computer Studies*, Vol. 43, pp.907-928.

Guarino, N., and P. Giaretta (1995). Ontologies and Knowledge Bases; Towards a Terminological Clarification, *Towards Very Large Knowledge Bases*, N.J.I. Mars (*ed.*), IOS Press, Amsterdam, The Netherlands.

Kowalski, R.A., (1989) The Treatment of Negation in Logic Programs for Representing Legislation, Proceedings of the Second International Conference on AI and Law, ACM Press, New York, USA.

Kant, I. (1781). *Critik der reinen Vernunft*, (Critique of Pure Reason), First edition, Riga 1781.

Kralingen, R.W. van (1995). Frame-based Conceptual Models of Statute Law, Computer/Law Series, Kluwer Law International, The Hague, The Netherlands.

McCarty, L.T. (1995). An Implementation of Eisner v. Macomber, *Proceedings of the Fifth International Conference on AI and Law*, ACM Press, New York, USA, pp.276-286.

Moles, R.N., and S. Dayal (1992). There is more to Life than Logic, *Journal of Information Science* (draft version), Vol. 3, No. 2, pp.188-218.

Routen, T.W. (1989). A Framework for Legal Knowledge Base Construction, *Proceedings of the Second International Conference on AI and Law*, ACM Press, New York, USA.

Searle, J.R. (1964). How to Derive "Is" from "Ought", *Philosophical Review*, Vol.73, pp.43-58.

Sergot, M.J., F. Sadri, R.A. Kowalski, F. Kriwaczek, P. Hammond, and H.T. Cory (1986). The British Nationality Act as a Logic Program, *Communications of the ACM*, Vol.29, No. 5, pp.370-386.

Sergot, M.J. (1991). The Representation of Law in Computer Programs, T.J.M. Bench-Capon (*ed.*), *Knowledge Based Systems and Legal Applications*, Academic Press, London, UK, pp.3-68.

Valente, A. (1995). *Legal Knowledge Engineering; A Modelling Approach*, University of Amsterdam, The Netherlands, IOS . Press, Amsterdam, The Netherlands.

Visser, P.R.S. (1995). *Knowledge Specification for Multiple Legal Tasks; A Case Study of the Interaction Problem in the Legal Domain*, Computer/Law Series, No. 17, Kluwer Law International, The Hague, The Netherlands.

Visser, P.R.S., and T.J.M. Bench-Capon (1996). On the Reusability of Ontologies in Knowledge System Design, *Seventh International Workshop on Database and Expert Systems Applications* (DEXA'96), Zurich, Switzerland, pp.256-261.

Visser, P.R.S. and T.J.M. Bench-Capon (1997). A Comparison of Four Ontologies for the Design of Legal Knowledge-Systems, *Artificial Intelligence and Law*, special issue after JURIX'96 (to appear).