

Legal Knowledge Based Systems
JURIX'96
Foundations of legal knowledge systems

The Foundation for Legal Knowledge Systems

Editors:
R.W. van Kralingen
H.J. van den Herik
J.E.J. Prins
M. Sergot
J. Zeleznikow

Deep models, ontologies and legal knowledge based systems
T.J.M. Bench-Capon and P.R.S. Visser

3-14

The publisher is aware that, in spite of all care taken by editors, authors and publisher, shortcomings may occur in this publication. For these, however, editors, authors and publisher cannot accept any liability.

Paper version: Tilburg University Press, ISBN 90-361-9657-4, NUGI 699

1996 JURIX The Foundation for Legal Knowledge Systems

Alle rechten voorbehouden. Niets uit deze uitgave mag worden verveelvoudigd, opgeslagen in een geautomatiseerd gegevensbestand, of openbaar gemaakt, in enige vorm of op enige wijze, hetzij elektronisch, mechanisch, door fotokopieën, opnamen, of op enige andere manier, zonder voorafgaande schriftelijke toestemming van de stichting JURIX.

Voor zover het maken van kopieën uit deze uitgave is toegestaan op grond van artikel 16b Auteurswet 1912 j° het Besluit van 20 juni 1974, Stb. 351, zoals gewijzigd bij het Besluit van 23 augustus 1985, Stb. 471 en artikel 17 Auteurswet 1912, dient men de daarvoor wettelijk verschuldigde vergoedingen te voldoen aan de Stichting Reprorecht (Postbus 882, 1180 AW Amstelveen). Voor het overnemen van gedeelte(n) uit deze uitgave in bloemlezingen, readers en andere compilatiewerken (artikel 16 Auteurswet 1912) dient men zich tot de stichting JURIX te wenden.

CONTENTS

Deep models, ontologies and legal knowledge based systems <i>T.J.M. Bench-Capon and P.R.S. Visser</i>	3
The formal specification of a legal ontology <i>Pepijn Visser and Trevor Bench-Capon</i>	15
An integrated view on rules and principles <i>Bart Verheij</i>	25
A generic model for the interpretation of vague norms <i>Jeannette Quast, Jaap van den Herik and Leo Aarts</i>	39
Salomon: automatic abstracting of legal cases for effective access to court decisions <i>Caroline Uyttendaele, Marie-Francine Moens, Jos Dumortier</i>	47
A task-based hyperindex for legal databases <i>Luuk Matthijssen</i>	59
A framework for self-explaining legal documents <i>L. Karl Branting and James C. Lester</i>	77
Word use in legal texts: statistical facts and practical applicability <i>Kees van Noortwijk and Richard V. De Mulder</i>	91
Automating legal reasoning in discretionary domains <i>Andrew Stranieri and John Zeleznikow</i>	101
Legal reasoning by structural analogy based on goal dependent abstraction <i>Tokuyasu Kakuta, Makoto Haraguchi and Yoshiaki Okubo</i>	111
Using set-of-support control strategy to deal with indeterminacy in legal reasoning <i>Bipin Indurkha</i>	123
Constructing Normative Rules <i>Nienke den Haan and Joost Breuker</i>	135

DEEP MODELS, ONTOLOGIES AND LEGAL KNOWLEDGE BASED SYSTEMS

T.J.M. Bench-Capon and P.R.S. Visser
Department of Computer Science, University of Liverpool
P.O. Box 147, L69 7ZF Liverpool, United Kingdom
email: {T.J.M.Bench-Capon, P.R.S.Visser}@csc.liv.ac.uk

Abstract

In this paper we explore the trend towards the production of “ontologies” as part of the development of knowledge based systems, both in general AI and in the legal domain in particular. We discuss four examples of this kind of work in the legal domain and identify areas on which future work might be directed.

1 Introduction

Since their introduction, knowledge based and expert systems have attracted a mixed reaction, mingling excitement at their potential with dissatisfaction with their various limitations. This has been true both in general AI, and in the legal field. Much work has therefore concentrated on understanding the limitations, and developing methods for building the systems which can mitigate the limitations. The overall trend has been away from seeing the process as one of encoding heuristics derived from an expert, towards modelling the domain on which the expertise operates. Further it has become recognised that modelling requires as a precondition that the domain be conceptualised, and that it is important that the assumptions incorporated in the conceptualisation be made explicit. Such an explicit conceptualisation has been called an ontology (Gruber 1991).

In this paper we will examine this trend. Following a brief overview of developments in general AI we will focus on the legal domain. In particular we will look at four different approaches to conceptualising the legal domain for the purpose of building legal knowledge based systems.

2 Knowledge Based Systems

The original conception of a knowledge based system was that the knowledge base would contain a set of empirical associations, such as might be formed by an expert in the course of practising in the domain, encoded in some executable formalism, typically production rules. MYCIN (Buchanan and Shortliffe 1985) remains the classic example.

Such first generation systems came under much criticism. See for example the papers in Bobrow (1984). We can point to four particular classes of criticisms:

- *Knowledge acquisition*: knowledge was difficult to acquire since experts often had difficulty in making their expertise explicit, and in stating their empirical associations with sufficient completeness.
- *Robustness*: the systems were brittle: they were unable to deal with situations not anticipated or overlooked by the expert.
- *Maintenance and Reuse*: the empirical associations had been formed by the expert in the context of the performance of a particular task, and so were difficult to adapt to other tasks, or to changes in the domain.
- *Explanation*: the systems gave poor quality explanations since the empirical associations did not reflect the underlying causal mechanisms of the domain.

2.1 Deep Models

Such systems were often characterised as “shallow” and an initial response to this was to advocate the use of “deep” models, which would explicitly model the various objects of the domain and the connections between them, and solve tasks not by applying acquired “rules of thumb”, but by reasoning about the model from first principles (Bobrow 1984). Such an approach was particularly suited to domains such as fault diagnosis. As an example consider the full adder used in Reiter (1987). In this problem we have a circuit which comprises two AND-gates, two XOR-gates and one OR-gate, connected in a particular way. The problem lends itself naturally to model based diagnosis. The gates can be modelled as having two inputs and an output, and the different kinds of gate by the functional relationship between their inputs and their outputs when they are working. The circuit can now be described by stipulating how the gates are connected together. A fault in the circuit is diagnosed by describing the initial inputs and the resulting outputs, and reasoning about which gates can be consistently believed to be working. Contrast this with the “shallow” approach, which would take certain patterns of input/output and associate them with certain faults. This domain is well suited to the model based approach, since it can be decomposed into components in a natural way. We can see how the four difficulties with rule based systems are addressed:

- *Knowledge acquisition* becomes a matter of identifying the components and describing the properties of the components. In this domain at least this is straightforward, as there is only one sensible way to do it.
- *Robustness* is achieved since situations can be simulated by the model, and so its coverage is determined by which situations were anticipated by the expert, but by the detail which the modelling includes.
- *Maintenance and Reuse* is helped because the components are readily separated and can be reconfigured to give other circuits. The independence and self-sufficiency of the components allow them to be used in other contexts. Moreover, we can use the circuit in other ways: for example either to deduce inputs on the basis of given outputs, or to deduce outputs on the basis of given inputs. We can even deduce “shallow” rules from the model if we wish.
- *Explanation* can now be couched in terms of the model in which the operation of the circuit is simulated, and the causal relations are preserved.

Deep models thus provided an answer to the problems of shallow systems - provided the domain was one which lent itself to being modelled in this way.

2.2 From Deep Models to Ontologies

So attractive was this approach, that it was felt desirable to develop a generally applicable methodology which would capture these benefits for any kind of knowledge based system. CommonKADS (Breuker and van der Velde 1994) can serve as an example. Key features of these methodologies were that they should enforce a clear separation between domain knowledge and control knowledge, and that the modelling of domain knowledge should be seen rather as encoding domain models rather than expertise. The difficulty in applying these methodologies, however, lay in the fact that whereas in the full adder domain there is an obvious decomposition of the circuit into gates and connections, in other domains the correct decomposition is far from obvious. Essentially, constructing the model depended on first conceptualising the domain. For principled development it is important that this conceptualisation be made explicit, and this gave rise to the notion of an ontology, a term popularised particularly by Gruber (1991). In some domains such as the VT domain (Gruber and Olsen 1996), the ontology looks very like the kind of deep model found in the full adder domain: in others the ontology remains at a considerably higher level: for examples

see Schreiber et al (1995). In either case, the production of an ontology as part of the domain model is now seen as an essential step in the principled development of knowledge based systems.

3 Legal Knowledge Based Systems

We can see a similar line of development in legal knowledge based systems. Initially such systems in the legal domain were developed like any other expert system, knowledge being elicited from an expert and encoded by a knowledge engineer. Excellent examples of this kind of system can be found in the Latent Damage Advisor of Capper and Susskind (1988), or the Nervous Shock Advisor of Smith and Deedman (1987). Although these systems were reasonably successful, they did indeed suffer from the drawbacks listed above: they required a considerable amount of effort from the expert; their coverage was restricted to situations envisaged by the expert; they were difficult to maintain and enhance and their explanations tended to be pieces of specially written text rather than recovered from the rules. Although systems of this style continued to be built - often prompted by experts in some particular field, there was a desire to find some more general approach to rectifying these defects.

One response to the problems of shallow systems was to base systems not on elicited expertise, but on an explicit representation of the applicable legislation. An early example is the representation of the British Nationality Act as a logic program (Sergot et al 1986). Advocates of this method claimed that it would provide many of the advantages of deep models (Bench-Capon 1989).

The obvious problem with the approach was that by restricting the knowledge base to legislation, limits were placed on the usefulness of such systems. Much of the expertise deployed in legal reasoning is directed towards the interpretation of legislation. Moreover, methods of formalisation were not as clear cut as might have been originally assumed. Subsequent developments of this approach concentrated on a variety of aspects:

- General methods for representing particular constructs not straightforwardly translated into predicate logic (*e.g.*, Routen and Bench-Capon 1991);
- Methods of integrating an explicit representation of legislation with interpretative expertise (*e.g.*, Bench-Capon 1991);
- Guidance for producing maintainable systems (*e.g.*, Bench-Capon and Coenen 1992).

Thus the development work was concentrated on either issues of how to engineer the systems better, or of how to integrate representations of legislation with expertise. Most radical criticisms of explicit formalisations of legislation, however, did not concern the engineering aspects at all, but rather the implicit conceptualisation of the legal domain that underlay this approach.

Most critics characterised the underlying conceptualisation as being that “law is a set of rules”. Moles (*e.g.*, 1992) and Leith (1990) are representative of these critics. Such a conceptualisation may be attributed to the approach, but is probably more fairly attributed to other rule-centric approaches such as Susskind (1987). In fact a truer characterisation of the conceptualisation used in the representation of legislation approach is that law is a set of definitions. If we look at the particular areas of law where this approach has been relatively successful, (the best examples being social security law) we see that much of the legislation is indeed concerned with defining concepts. In so far as the legislation to be represented is a set of definitions, then explicitly stating the necessary and sufficient conditions for the application of a concept and using deduction to see whether a concept is applicable is a viable conceptualisation, and one which has clear computational merits. At least it can be seen to have a role in the process of reasoning about this kind of law (Bench-Capon 1994).

Once we have made explicit the conceptualisation in this approach we can see its limitations; both in the nature of the law to which it can be applied, which must be

largely definitional in nature, and in the use that we can make of the system, in that it cannot take us beyond the logical application of definitions. In some cases - the leading example is the application of social security regulations by lay adjudicators, - such a system may be what is wanted. Here the chief concern is for consistent, impartial decisions, and adversarial reasoning is largely absent. Understanding the nature of the conceptualisation remains, however, vital if the approach is to be applied appropriately, so that we can see where it applies and where it does not. The limitations make it clear that the approach will not necessarily be appropriate for other areas, and if we are to represent these we must remedy the defects of first generation knowledge based systems in other ways.

4 Deep Conceptual Models in the Legal Domain

In this section we discuss two critics of the approach based on an explicit representation of legislation.

4.1 McCarty's Language for Legal Discourse (LLD)

One of the first to advocate the use of deep models in the legal domain was Thorne McCarty. In McCarty (1989) he wrote, referring to two earlier papers:

“I advocated the development of “deep conceptual models”. My motivation was both practical and theoretical. On the practical side, I argued that our long term goal should be an integrated analysis/planning/retrieval system that matches as closely as possible the way a lawyer actually thinks... On the theoretical side my work with Sridharan on the TAXMAN project had clarified the importance of a domain model in any attempt to model the arguments of lawyers in hard cases. For both purposes, I claimed, deep conceptual models are essential”

He then went on to attempt to clarify what he meant by “deep conceptual model”. Having noted that there were a number of common sense categories underlying the representation of a legal problem domain, including such things as space, time, causality, permission, obligation, intention knowledge and belief, he said:

“The idea is is to select a small set of these common sense categories.. and then develop a knowledge representation language that faithfully mirrors the structure of the set... If a language of this sort could be developed it would provide a uniform framework for the construction of a legal analysis/planning/retrieval system, and a solid foundation for for further theoretical work.”

McCarty's earlier work on permissions and obligations (McCarty 1983) could be seen as an effort to explicate two of these categories, and the paper quoted above introduced his Language for Legal Discourse, which was intended as a first step in the more ambitious aim set out in the second quotation.

Unfortunately a language on its own is not enough: although this can provide a framework, the proof of its effectiveness can only given by an evaluation of the domain models constructed with it. These are, however, few and far between. McCarty himself, has attempted to represent the concept of ownership, but this has not been widely published (McCarty 1993), and the best flavour of the model that would be produced using the framework is found in McCarty (1995). In that paper he explicitly criticises the conceptualisation underlying the “law as definitions” approach (and his own TAXMAN I):

- “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.”

His alternative conceptualisation identifies three components to a legal concept: a set (possibly empty) of invariant necessary conditions; a set of exemplars providing sufficient conditions; and a set of transformations expressing the relationships among the exemplars. As a result:

“The application of a concept to a new factual situation automatically modifies the concept itself”

presumably by extending the set of exemplars, or the set of transformations. The rest of the paper goes on to illustrate this conceptualisation in the context of the famous *Eisner vs Macomber* case used in all the TAXMAN experiments.

The conceptualisation has two interesting features. First it is specifically directed towards computational solutions in the domain, although in the absence of a full implementation, its success in this respect must be taken on trust. Second, it puts specific emphasis on the dynamics of legal concepts, and of the role played by argument in realising that dynamic nature. As such it is particularly suited to the kind of hard cases that might receive the attention of the Supreme Court - which has always been McCarty’s interest. On the other hand if the case is easy - already covered by the necessary conditions or some exemplar, then it seems to reduce to much the same kind of thing as the definitional conception, with the exception that because the sufficient conditions are exemplars, we need to understand some appropriate matching algorithm to determine whether or not the sufficient conditions so represented are applicable or not. Many questions, however, remain to be answered:

1. To what extent does the LLD capture the commonsense notions of, for example, permission and obligation? Certainly there is no universal consensus that McCarty’s version is successful - for example Jones has raised several objections (Jones 1987).
2. How reusable are the concepts McCarty defines? For example ownership, although designed in McCarty (1993) for ownership of stocks, needs to be capable of application to land, goods, money, rights and so on, if it is to fulfil its part in a deep conceptual model. This can be determined only if the definition is used in other applications.
3. Is the conceptualisation really appropriate for all law, or only hard cases, and the kind of hard cases that McCarty has looked at? Again this can be resolved only with wider use.
4. How important is it that the sufficient conditions are held as exemplars? At this point McCarty’s chosen problem solving method - prototypes and deformations - seems to be determining the representation. On the other hand it does link this conceptualisation to that which underlies the Case Based Reasoning approach to legal systems, exemplified by HYPO (Ashley 1990) and others.
5. How appropriate is the conceptualisation at all? McCarty’s justification lies in an appeal to its relation of the three points listed above which are “familiar to all lawyers”, coupled with a suggestion that it conforms to the requirements of Levi’s account of legal reasoning (Levi 1942). It would, however, be instructive to see to what extent it stands up to jurisprudential criticism.

This work - which represents one of the few attempts to apply the idea of a deep model to the legal domain - is worthy of more attention than it has so far received. Until it receives such attention, however, it is difficult to evaluate it, since as it purports to be a framework for implementing many applications, and for conducting theoretical research, and its utility as such cannot be demonstrated until it is used.

4.2 *Stamper's NORMA*

Another radical critic of the approach which relies on an explicit representation of legislation is Ronald Stamper. Stamper has criticised the use of traditional logics for the representation of (legal) knowledge because they suffer from some important semantic problems (Stamper, 1991). Briefly stated, traditional logics rely on symbolic representations that have only a very weak connection to the real-world concepts they intend to denote. In particular, symbolic representations rely (according to Stamper - invalidly) on notions such as truth, individuality, and identity. Accordingly, expressing legal knowledge in the form of rules is an over simplification of what legal knowledge is about. To overcome these problems Stamper argues that there is need to escape from the frame of reference within which the classical logic is created. Building on his LEGOL work (see Stamper, 1980), he proposed the NORMA formalism (Stamper, 1991). NORMA, which means "logic of norms and affordances", is based on two main philosophical assumptions: (1) there is no knowledge without a knower, and (2) the knowledge of a knower depends on its behaviour (Stamper, 1996). Using NORMA the entities in the world are described by their behaviour rather than by assigning them an individuality and truth values. The main ontological concepts are:

- (a) *agents*, organisms standing at the centre of reality. They gain knowledge, regulate and modify the world by means of actions. Agents are responsible for their actions. The concept of an agent can be extended to include groups, teams, companies, social agents or nation states even.
- (b) *behavioural invariants*. One of the underlying ideas of NORMA is that the knowledge of the world is to be described by features that remain invariant over some time and that these features are found in the behavioural characteristics objects. For instance, a cup is described by the ability to hold liquids, the noise it makes in hitting various surfaces, the visual shape it displays etc. To express this, NORMA uses the construct of a behavioural invariant. A behavioural invariant is a description (*e.g.*, using verbs, nouns, or adjectives) of a situation whose features remain invariant.
- (c) *realisations*. Agents realise situations by performing actions. The realisation of a situation - a realisation - is specified as the combination of (1) an agent and (2) a behavioural invariant, shortly written as Ax (the situation, denoted by behavioural invariant x , that is realised by agent A). An example of a realisation Ax is John walks. Different kinds of realisations are recognised, for instance, Ax^* (denoting the ability of A to realise x), $Ax@$ (denoting the authority of A to realise x), $Ax+$ (denoting A starts to realise x), $Ax-$ (denoting A finishes the realisation of x), and $Ax\#$ (denotes that x can be divided into individuals, cm. classes and objects).

By combining behavioural invariants composite realisations can be made. We here mention the most important composite realisations: Axy (denoting that A cannot realise y without first realising x), $A.x.y$ (denoting that x is a part of A and y is a part of x), $A(x \text{ while } y)$, $A(x \text{ or while } y)$, $A(x \text{ while not } y)$, $A(x \text{ whenever } y)$ (denotes a norm; x should be realised whenever y is realised); $A(x \text{ then } y)$ (denotes a norm, if x is realised then y should be realised), A "Ax" (denoting that an agent A can state another agent B to bring about x , for instance by commanding or suggesting), and $A(a:b:c)d$ (denoting that a , b , and c are instances of d). For a more extensive discussion of these notions see Stamper (1991; 1996).

In a sense we can see the NORMA work as attempting to fulfil the same role as McCarty's LLD: providing a general representational language for law which can serve as a framework to construct models of particular domains. That this approach represents a novel conceptualisation, and a challengingly different one to those found elsewhere in AI and Law, cannot be doubted. Its utility is not, however, as yet proven. First, there are no substantial representations of law using this system in Stamper's published work: there are some fragments which give some of the flavour of the representation, but the material is far from complete. Second the definition of the formalism is also rather sketchy: it is unclear for example whether the different kinds of realisation are intended to be exhaustive.

Next we need to ask questions about NORMA with respect to legal theory. Its basis lies rather in semiotics than legal theory, and is intended to focus on social domains, so it is unclear whether it is intended to embody a legal theory or to provide the means for expressing different legal theories. Finally, the computational implications of this conceptualisation remain an open question.

In order to come to a good assessment of the worth of this conceptualisation again work is needed to provide a substantial body of material employing it, and then to answer the questions posed above.

5 Ontologies in the Legal Domain

The last two conceptualisations that we will discuss explicitly proclaim themselves as ontologies, the functional ontology of Valente (1995), and the frame based ontology of van Kralingen (1995) and Visser (1995). In both of these cases the provision of an explicit, formally described, conceptualisation of the legal domain is an avowed aim of their work.

5.1 Valente's Functional Ontology

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, determined by social goals. Its main function is reacting to social behaviour. This main function can be decomposed into six primitive functions, each corresponding to a category of primitive legal knowledge.

- (1) *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.
- (2) *world knowledge*: which describes the world that is being regulated in legal terminology, and can be considered as an interface between common sense knowledge and normative knowledge.
- (3) *responsibility knowledge*: this is the knowledge which enables responsibility for the violation of norms to be ascribed to particular agents.
- (4) *reactive knowledge*: which describes the sanctions that can be taken against those who are responsible for the violation of norms
- (5) *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.
- (6) *creative knowledge*: which states how items of legal knowledge can come into being and cease to be.

Valente's ontology is described in Gruber's ontology language Ontolingua (Gruber 1992). It has also formed the basis of a system ON-LINE 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 knowledge base, 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. This is rather different from the role of the conceptualisation in McCarty and Stamper, where it is supposed to form the basis of a general representation language embodying some fundamental legal notions.

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 is relatively complete, the domain model constructed within this framework for the ON-LINE system, seems 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.”

Thirdly one can again ask about the relationship between this ontology and legal theory. Valente argued that ontologies should provide the “missing link” between legal theory and legal information systems, but we are not jurisprudentially competent to discuss whether his functional ontology does so.

5.2 The Frame Based Ontology of Van Kralingen and Visser

The final conceptualisation we will discuss was motivated by a desire to improve development techniques for legal knowledge based systems, and in particular a desire to reduce the task dependency of knowledge specifications. This ontology distinguishes between an ontology which is intended to be generic to all law, and a statute specific ontology which is required to supplement it in any particular case.

An Ontolingua version of the generic ontology appears in this volume (Visser and Bench-Capon 1996) and so it will not be described here. It has formed the basis for two applications (FRAMER) in Dutch Unemployment Benefit Law (Visser 1995b). Visser (1995) provides one of the very few examples of an ontology pursued through to a full implementation.

Probably it is the conceptualisation of the domain found in the generic ontology that is most comparable to the other conceptualisations we have so far discussed. The usual questions may be asked about this: are the three entities identified sufficient to represent a variety of legal domains, and are the attributes assigned to those entities necessary and exhaustive? Turning to the statute specific ontology we need to ask whether the method used to produce it would be applicable across a variety of different types of legislation. It was found to be a helpful basis for development in Unemployment Law, but this is very definitional in nature.

6 Discussion

What the above is intended to show is that the first generation idea of encoding expertise is likely to give rise to a system which is inadequate in terms of knowledge acquisition, maintainability, reuse, robustness, and the ability to provide satisfactory explanations. The deep model answer, however, is not immediately applicable to law: the identification of components in law - well agreed in areas such as fault diagnosis - is rather problematic in our domain. Advocates of deep models have tended to produce a language for representing legal domains, in the fashion of Stamper and McCarty, or more generally a formal, explicit conceptualisation of the domain, such has now become commonly called an ontology. The term is openly embraced by two of the approaches discussed above, and Stamper describes some of the representation produced in NORMA as "ontological charts". Nor do we think that McCarty would reject the description of LLD as an ontology. Such conceptualisation is an essential preliminary step if any deep modelling is to be done.

What is common to all these approaches - different as they are - is the belief that in order to make progress it is necessary explicitly to state the assumptions that are being made about the domain - to make explicit the conceptualisation that underlies the systems that may be built. The potential gains can be discovered by looking again at the approach based on an explicit representation of legislation. This was never an explicit ontology, but there was an implicit ontology - the conceptualisation of the legal domain as a set of definitions which could be represented as a set of first order predicate calculus definitions from which the consequences could be deduced. However inadequate this may be as a general characterisation of law, it has attracted a great deal of attention in the AI and Law community. Systems have been built, and criticised, and refined, and legal theorists have railed against its limitations. Many people, from different perspectives and different background have worked on this conceptualisation with different motivations. As a result, what can be done with this conceptualisation is now pretty well understood. We know:

- its limitations in the sense of the sort of legislation to which it can be applied;
- its limitations in the sense of the sort of system that will result from its application;
- how to build and use legal information systems within this conceptualisation.

The four conceptualisations discussed in this paper, however, have not received anything like the same attention; for the most part their use has been confined to a single author or group. The result is that while we can examine these conceptualisations, and identify their various attractions, we are left with questions which cannot be resolved until a body of work has been produced, which can form the basis for a reliable evaluation. What is needed is a programme of work which will include the following activities:

- (1) The explicit expression of conceptualisations in a complete and rigorous form. Ideally the expression of different conceptualisations should be readily comparable. The trend to use a common ontology language such as Ontolingua to express ontologies is very welcome.
- (2) The conceptualisations need to be used across a variety of different applications. In so far as the ontology is meant to be generic, it is important to establish - as was done for the "law as definitions" conceptualisation - where it can and cannot be effectively applied.
- (3) The use of different conceptualisations of the same area of law. This is important if we are to get a real understanding of the difference that different conceptualisations make.

- (4) Discussion of the conceptualisations from a jurisprudential standpoint. At present the conceptualisations are predominately produced by those whose main interest is computational. The role of legal theory in this enterprise is to uncover the viability of these conceptualisations from the point of view of legal theory.

This is a programme of work which is too much for a single researcher, or a single group. Moreover the work needs to be done from a plurality of perspectives. Hopes for progress in AI and Law require that the community works in a more synergistic way, explicitly building on and extending the work of others.

References

- Ashley, K.D., (1990). *Modeling Legal Argument*, MIT Press: Cambridge, Mass.
- Bench-Capon, T.J.M., (1989) Deep Models, Normative Reasoning and Legal Expert Systems, *Proceedings of 2nd International Conference on AI and the Law*, Vancouver, 1989, ACM Press, pp 37-45.
- Bench-Capon, T.J.M., (1991). Practical Legal Expert Systems: the Relation Between a Formalisation of Law and Expert Knowledge. In Bennun and Narayanan (eds) *Computers, Law and AI*, Ablex, 1991, pp191-201.
- Bench-Capon, T.J.M., and Coenen, F.P., Isomorphism and Legal Knowledge Based Systems, *AI and Law*, Vol 1, No 1, pp65-86.
- Bench-Capon, T.J.M., (1994). Legal Theory and Legal KBS: A Computer Scientist's Perspective, in Prakken, H., Muntjewerff, A.J., Soetman, A., (eds) *Legal Knowledge Based Systems: The Relation with Legal Theory (JURIX 1994)*, Koninklijke Vermaede, Lelystad, pp 33-42.
- Bobrow, D., (ed) (1984) *Qualitative Reasoning About Physical Systems*, published as *Artificial Intelligence*, vol 24.
- 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.
- Buchanan, B.G., and E.H. Shortliffe (1985). *Rule-Based Expert Systems; The MYCIN experiments of the Stanford Heuristic Programming Project*, Addison-Wesley Publishing Company, Massachusetts, United States.
- Capper, P.N., and Susskind, R.E., (1988). *Latent Damage Law - The Expert System*, Butterworths, London.
- Gruber, T.R. (1991). The Role of Common Ontology in Achieving Sharable, Reusable Knowledge Bases, *Principles of Knowledge Representation and Reasoning, KR'91, Proceedings of the Second International Conference*, J. Allen, R. Fikes, and E. Sandewall (eds.), pp.601-602, Cambridge, Massachusetts, United States.
- Gruber, T.R. (1992). *Ontolingua: A Mechanism to Support Portable Ontologies*, Knowledge Systems Laboratory, Stanford University, Stanford, United States.
- Gruber, T.R., and G.R. Olsen (1996). The Configuration Design Ontologies and the VT Elevator Domain Theory, *Int. Journal of Human-Computer Studies*, Academic Press, Vol.44, pp.569-598.
- Jones, A.I.J., (1987) On the Relationship Between Permission and Obligation in *Proceedings of the Second International Conference of AI and Law*, ACM Press: New York, pp 164-169.
- Kralingen, R.W. van (1995). *Frame-based Conceptual Models of Statute Law*, Computer/Law Series, No. 16, Kluwer Law International, The Hague, The Netherlands.
- Leith, P., (1990) *Formalism in AI and Computer Science* Ellis Horwood: Chichester.
- Levi, E.H., (1942) *An Introduction to Legal Reasoning* University of Chicago Press.
- McCarty, L. Thorne, (1983) Permissions and Obligations in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pp 287-294.

- McCarty, L. Thorne, (1989). A Language for Legal Discourse I. Basic Features In *Proceedings of the Second International Conference of AI and Law*, ACM Press: New York.
- McCarty, L. Thorne, (1993) OWNERSHIP: A Case Study in the Representation of Legal Concepts, Presented at a *Conference in Celebration of the 25th Anniversary of the Istituto Documentazione Giuridica*, Florence, Italy, 1993, distributed in typescript.
- McCarty, L. Thorne, (1995) An Implementation of Eisner v. Macomber in *Proceedings of the Fifth International Conference on AI and Law*, AVM Press: New York, pp 276-286.
- Moles, R.N., (1992) Expert Systems - The Need For Theory in Gutters, C.A.F.M., et al (eds), *Legal Knowledge Based Systems and Information Technology in Law, JURIX 92*, Koninklijke Vermande, Lelystad, pp 113-122.
- Reiter, R. (1987). A Theory of Diagnosis From First Principles, *Artificial Intelligence* 32, pp. 57-95.
- Routen, T., and Bench-Capon, T., (1991). Hierarchical Formalisations *International Journal of Man Machine Studies* 35, pp 69-93
- Schreiber, A.T., Wielinga, B.J., and Jansweijer, W.H.J., (1995). *The KACTUS view on the 'O' word*. Technical Report, ESPRIT Project 8145, KACTUS, University of Amsterdam
- Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., Cory, H.T. (1986). The British Nationality Act as a logic program. *Communications of the ACM* 29, 5, pp 370-386.
- Smith, J.C., and Deedman, C. (1987). The Application of Expert Systems Technology to Case Based Law, in *Proceedings of the First International Conference on AI and Law*, ACM Press: New York.
- Stamper, R.K. (1980). LEGOL: Modelling Legal Rules by Computer, *Computer Science and Law*, Niblett, B., (ed.), Cambridge University Press, Cambridge, UK.
- Stamper, R.K. (1991). The Role of Semantics in Legal Expert Systems and Legal Reasoning, *Ratio Juris*, Vol. 4, No. 2, pp.219-244.
- Stamper, R.K. (1996). Signs, Information, Norms and Systems, *Signs at Work*, B. Holmqvist, and P.B. Andersen (eds.), De Bruyter, Berlin, Germany.
- Susskind, R.E., (1987) *Expert Systems in Law*, Clarendon Press: Oxford.
- Valente, A. (1995). *Legal Knowledge Engineering; A Modelling Approach*, 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. (1995b). *FRAMER: Source Code of a Legal Knowledge System that performs Assessment and Planning*, Reports on Technical Research in Law, University of Leiden, Leiden, the Netherlands, Vol. 2, No 1.
- Visser, P.R.S., and T.J.M. Bench-Capon (1996). The Formal Specification of a Legal Ontology, *Proceedings of the Ninth International Conference on Legal Knowledge Based Systems, JURIX-96*, this volume.