

Legal knowledge based systems  
JURIX 94  
The Relation with Legal Theory

The Foundation for Legal Knowledge Systems

Editors:

H. Prakken  
A.J. Muntjewerff  
A. Soeteman

P.W. Brouwer, Legal Knowledge Representation in the Perspective of Legal Theory, in:  
H. Prakken, A.J. Muntjewerff, A. Soeteman (eds.), Legal knowledge based systems  
JURIX 94: The Foundation for Legal Knowledge Systems, Lelystad: Koninklijke  
Vermande, 1994, pp. 9-18, ISBN 90 5458 190 5.

More information about the JURIX foundation and its activities can be obtained by  
contacting the JURIX secretariat:



Mr. C.N.J. de Vey Mestdagh  
University of Groningen, Faculty of Law  
Oude Kijk in 't Jatstraat 26  
P.O. Box 716  
9700 AS Groningen  
Tel: +31 50 3635790/5433  
Fax: +31 50 3635603  
Email: [sesam@rechten.rug.nl](mailto:sesam@rechten.rug.nl)

# Legal Knowledge Representation in the Perspective of Legal Theory

P.W. Brouwer  
*University of Amsterdam*

## 1 Introduction

One of the main problems in the field of Artificial Intelligence and Law, perhaps the main problem [McCarty, 1990], is the construction of adequate representations of legal knowledge. The notion of adequacy involved has at least two aspects. Adequacy is partly determined by typical AI-considerations like computational tractability and the efficient maintenance of legal knowledge based systems. On the other hand, no representation of law is adequate if it does not mirror the law. The representation of law has to be equivalent, in some sense, to the law that is represented. If one is satisfied with anything less, there will be no faithful, no valid representation. What the system offers, is at best an approximation, and at worst a distorted picture of law. In this paper I will argue that (analytical) legal theory may be of some help in tackling the problem of valid representations, since it addresses itself to the kind of questions that will have to be solved in order to create such representations of law. Given the present state of the art however, one should not expect legal theory to offer ready-made solutions.

The development of AI and Law has shown an increasing awareness of problems of a legal theoretical nature. But it has been mainly the theory of legal reasoning that has attracted attention. Careful and thorough examinations of the nature of legal knowledge are relatively rare. Inquiries into these aspects of legal knowledge representation may counterbalance the somewhat one-sided emphasis on legal reasoning. In the following I will focus therefore on problems of legal knowledge representation. I will not, however, discuss the pros and cons of AI representation techniques like logic programming, frames, rules and semantic nets. It is not the competence of a legal theorist as such to tell a knowledge engineer what techniques he ought to use. It is to show him the characteristics of the legal domain, so that he may adjust his techniques to what that domain requires.

## 2 Problems of legal knowledge representation

To begin with I will list some problems which I consider to be basic problems of legal knowledge representation from a legal theoretical point of view. These problems, that are highly intertwined, concern:

1. the nature of law
2. the identification of the categories of legal knowledge
3. the components and internal structure of each of these categories

#### 4. the relations between different categories

To a certain extent, these problems cannot be evaded [Gold & Susskind, 1985, p. 307]. In constructing a legal knowledge based system, a solution to each of them is necessarily assumed. So the question is not whether a legal theory is useful or not. It is there, whether you like it or not. That, of course, does not imply that sufficient attention has been paid to the problems involved. The solutions may have been accepted more or less implicitly, and may be very rudimentary.

Any set of such solutions (assumptions) may be called a picture of law. An example of a set of this sort is the following: law consists of rules (solution to 1), these rules impose duties or confer powers (solution to 2), all rules have a conditional structure and consist of an act-description and a normative modality (solution to 3), and power-conferring rules make it possible that duty-imposing and power-conferring rules are created (solution to 4).

In view of the fact that assumptions of this sort are inevitable, it is obvious that legal theory, in as far as it addresses itself to the underlying problems, is of importance to research in the area of legal knowledge representation. Unless, of course, legal theory has nothing better to offer than what already has been adopted by legal knowledge engineers. I will argue that the present state of legal knowledge based systems does not allow that conclusion. Legal theory may still offer valuable insights. It enables us to judge the quality of the pictures of law that are the basis of these systems. It also may counteract a tendency in legal knowledge representation to underestimate the complexity of the legal domain. Many representations of law in fact are rather simple working models, that do not capture the subtleties of legal language. They often do not attain the level of accuracy that legal theory makes possible. In the following, I will provide some examples.

### 3 The nature of law

The most fundamental issue legal knowledge representation is confronted with, is the nature of law. What is it that has to be represented? What is law? It is a well-known fact that legal theory offers many different answers to that question. There is a variety of conceptions of law. So the question rises which conception has to be chosen for the purposes of legal knowledge representation and on what grounds that choice is to be made. My suggestion is, that these grounds are in principle no others than those that are generally adduced. Roughly speaking, it may be said that a conception may be chosen because it offers an adequate description of what is considered to be law (in a given community), because its adoption may enhance the quality of law, or because it does both.

Now some theories of law offer better perspectives for building legal knowledge based systems than others. And some other theories offer no perspective at all, at least according to their author. Dworkin for instance, believes that law (i.e. law as integrity) cannot be mechanized:

“No electronic magician could design from my arguments a computer program that would supply a verdict everyone would accept once the facts of

the case and the text of all past statutes and judicial decisions were put at the computers disposal” [Dworkin, 1986, p. 412]

Indeed, if law is not “exhausted by any catalogue of rules or principles” [Dworkin, 1986, p. 413], but seen as the enterprise of constructing increasingly better justifications for legal practice, a knowledge based system is practically impossible. No part of law, be it a statutory rule or a precedent, has a fixed meaning (the meaning, for instance, as currently assigned in legal practice). It derives its meaning from the best constructive interpretation, that depends on moral considerations and on the particularities of the cases to be decided, particularities that cannot be anticipated. So, according to this view, we may represent and formalize current interpretations and yet fail to capture the law. Tina Smith recently argued that this does not imply that a mechanization of legal reasoning is impossible. It would only imply that there is no certainty that the expert system’s solution is always right [Smith, 1994, p. 96]. In Dworkin’s conception however, a solution that is not right cannot be law. So if his conception offers an adequate description of law, one cannot succeed in building legal expert systems that rest on valid representations of law.

Assumptions concerning the nature of law not only determine the answer to the question whether faithful representations of law are possible or not. They also determine the object of representation. An example is Bench-Capon’s recent claim that many legal regulations can be represented as ‘tariffs that enable rational behaviour’ [Bench-Capon, 1994]. The basic idea is the following. Suppose we have two regulations. One of them states that A has a duty not to do actions of type B. The second states that if A does an action of type B, he will be liable to a penalty C. It is further assumed that an action of type B does not in itself constitute a moral wrong. It does not belong to the category of the *mala in se*, as some legal theorists would say. Now Bench-Capon claims that, on this condition, the informational content of the two provisions is captured by the second. ‘The regulations are not norms but rather serve to provide information as to the consequences of certain actions’. This information enables A to decide whether it is profitable or not to refrain from doing B. It offers him knowledge concerning ‘the price’ of actions of type B. It may be the case that a cost–benefit analysis shows him that doing B will maximise his personal utility, even if he has to pay the price (the sanction being inflicted). Being a rational person, he must then do B. If, on the other hand, he wishes to evade the sanction at all costs, he must avoid doing B.

It is clear what the main advantage of this approach would be. If the procedure works, it enables us to represent a considerable part of the legal domain without the use of deontic modalities. It diminishes the need for a deontic logic. As a consequence, the many unsolved difficulties of systems of deontic logic may – to a large extent – be circumvented.

Someone working in the field of legal theory will immediately recognize the gist of Bench-Capon’s proposal. There have been many attempts to recast duty–imposing norms as something else, and this is one of them. It is strikingly similar to Holmes’ treatment of legal duties:

‘Take again a notion which as popularly understood is the widest conception which the law contains – the notion of legal duty [...] We fill the word with all the content which we draw from morals. But what does it mean to a bad

man? Mainly, and in the first place, a prophecy that if he does certain things he will be subjected to disagreeable consequences by way of imprisonment or compulsory payment of money' [Holmes, 1992]

What to think of this proposal? From a legal theoretical point of view, the answer is obvious. The debate in legal theory has shown that such reductions are feasible, but of limited significance. There is no reason to deny that from an external point of view law *can* be seen as a system of tariffs that enable rational behaviour. The corresponding attitude towards law is characteristic for homo economicus, the calculating individual (not necessarily always 'a bad man'), who does not care about the legal norms as such. What he cares about are the unpleasant consequences of the violation of these norms (and in the case of 'positive' sanctions the pleasant consequences of obeying them). There is nothing wrong with the external point of view as such. If adequate, descriptions of law from this point of view provide knowledge about law.

From the internal point of view, however, it is a total failure, because it cannot explain in what way law is normative. It is of no use for those who participate in the practice of legal reasoning, and accept and use duty-imposing norms as *reasons* for their decisions and their conduct. From the internal point of view a duty constitutes a reason for action, irrespective of the fact whether a sanction may be inflicted in case of disobedience. Sanctions only provide auxiliary reasons for acting in accordance with legal norms. If we eliminate those norms from our representation of a legal domain, we cannot expect to give a faithful representation of law as seen from the internal point of view.

The claim, therefore, that the deontic modalities may be discarded because one does not need them for representing law, is too bold, even if that claim is restricted and does not cover the legal duties that pertain to (im)moral conduct. It has to be qualified further. Otherwise the internal point of view is simply defined out of existence [Hart, 1994, p. 91].

#### **4 Legal knowledge categories and the structure of law**

Legal knowledge is not homogeneous. The sources of law (statutes for instance, or precedents and treaties) contain thought-formations of many different types. In it figure the kernels [vonWright, 1963, p. 70] of duty-imposing, permissive and power-conferring norms, exceptions, definitions of legal concepts, provisions concerning the local or temporal range of application of norms, descriptions of offences, ordered pairs of case-descriptions and decisions, legal fictions, presumptions, criteria for legal causation, laws that formulate conditions for liability to sanctions, goals, values, principles, and priority rules. The list is not exhaustive. I will call these thought-formations fragments of law.

Now we may say that each of these types constitutes a category of legal knowledge. The advantage of that approach is that it impresses to one the diversity of legal thought-formations. The representation of causal knowledge, for instance, may require other techniques than the representation of norms or legal definitions. Such matters can only be decided after an examination of the characteristics of these types of knowledge. Part of that work has been done in legal theory.

Although it is important therefore to examine the features of the various fragments of law, in the first instance I would like to focus on something else. There is a presupposition, shared by many in legal theory, that behind this diversity of the legal sources more fundamental legal categories may be discerned. The reason for that assumption is the fact that most fragments of law do not have an independent status. They cannot be understood in isolation, because they derive their meaning from the position they have in a more extensive thought-formation, which is assumed to have a (relatively) independent status. There is law behind law. Without going into details, it may be said that this other picture of law, which is a picture of its deep structure, represents law as a set of complete norms for human behaviour. The basic idea, then is simple. Every fragment of law is part of something else: a complete norm.

Let me give an example. Article 53 of the Amsterdam APV (bye-law) prohibits (among other things) harassing other people in public buildings. The regulation does not contain a complete norm, since it lacks some essential ingredients. If the aim is to give prescriptions for conduct (and that is certainly one of the aims of article 53), a minimum requirement for successful regulation is that one provides the necessary information for those whose conduct one wants to regulate. Now both common sense and legal experience show, that a prescription ought to contain at least five components: an act-description, a description of the persons to whom the prescription is addressed, a deontic modality (obligatory or prohibited), an indication of time, and an indication of place [Brouwer, 1990]. Application conditions, that often are contained in one or more of the other components (but not in the deontic modality, of course), may be a sixth component. Together, the components constitute an abstract model of a complete norm (of the 'ought-to-do-type').

If we apply this model, it becomes clear that article 53 is not self-contained. It only gives us the kernel of a prescriptive norm. It provides information about the type of action that is forbidden and about the sort of places where those actions are not to occur (public buildings). It is silent however, with regard to the category of persons addressed, the stretch of time during which these actions are forbidden, and the location of the 'public buildings' involved.

If we read article 53 in isolation from its legal context, the most plausible reading (on the basis of general, non-legal presuppositions) would be, I suppose, that no person may ever harass other people in public buildings. The prohibition addresses everyone and covers all points in time. That however, is not what article 53 means. The general presuppositions have been replaced by special presuppositions that are to be derived from the context of this article. In the case at hand there are two domain-dependent presuppositions. The first is that the forbidden actions can only be actions (of persons who find themselves) within the territory of the municipality of Amsterdam. This presupposition rests on a power-conferring norm, that restricts the local power to legislate (in principle) to the territory of that municipality. The second specific presupposition is that the temporal range of application of article 53 is restricted to all points of time within the period of its legal validity (note that this would be different, if it were a retroactive law). Both presuppositions are essential for an adequate representation.

The idea of (relatively) independent and self-contained norms is extremely important in legal theory, because it is central in explicating the contents and the structure of law. The meaning of fragments of law can only be ascertained on the basis of their relations

to other fragments. What relations one assumes to exist, depends on a conception of what it is that links fragments to other fragments. The idea of complete norms (of some sort) is such a conception.

Although the importance of the distinction between independent norms and fragments of law (that are sometimes called 'dependent norms') is generally acknowledged in legal theory, there is no agreement about the sort(s) of norm that may be regarded as a complete and independent norm. Kelsen, who endorsed the distinction, once defended the view that there is only one kind of independent norm: the norm that prescribes or permits an 'act of coercion' (in most cases a sanction) under certain conditions [Kelsen, 1960, pp. 52–59]. All other fragments of law are 'dependent norms', that express some of these conditions. On this view, duties, rights, permissions and powers can all be recasted as fragments of independent norms. This extreme theory has been criticised, among others by [Hart, 1994, pp. 38–42]. It obscures the differences between diverse types of legal norms. It cannot do justice to the function and meaning of these norms in legal thinking.

The controversies in legal theory with regard to these fundamental legal knowledge categories, raises the question what guiding requirements control their identification. At least the following two can be mentioned. First, the set of these categories shall be comprehensive. It shall enable us to map an arbitrary fragment of law onto one or more complete norms. Secondly, it shall be possible to do this without changing or distorting the contents of the law. Legally accepted meanings should be preserved.

Now according to this view, legal theory builds models that enable us to map the fragments of law onto complete norms. It should be stressed, that the purpose of this mapping is not to recommend these models as the ideal form to express the contents of law. There are good reasons for current practice, in which an extensive use is made of fragments of law. On the whole, this technique avoids redundancy and facilitates the changes that may have to be made in a body of law. The point of complete norms is something else. They model the connections between various parts of the legal system and show us how the interaction between these parts determines their meaning. The construction of these models is a first step towards a semantical theory for legal language, which is indispensable if we want to achieve a precise account of the law that we would like to represent and formalize in legal knowledge based systems.

## **5 Abstract models and their theories**

If we strive for precision, it is not enough to have some relatively simple models, that hardly contain anything more than a list of the essential components of norms. Before embarking on the project of formalizing norms, we need a theory of those models. Basically, the function of such theories is twofold. In the first place, it should provide us with an adequate semantical analysis of the inner structure of the kind of norms concerned. Secondly, it should elucidate the concepts the model is founded on, such as 'action', 'time' and 'ought', as well as the legal concepts that are characteristic of the norms that are instantiations of that model, such as 'intentionally', 'cause' and 'liability'.

Let us consider some examples. The first of these concerns the inner structure of prescriptive norms. The question is, how the components of these norms are connected to each other. Now legal language, as it is used in the sources of law, shows that the location

of a component within the norm–formulation as a whole may vary. The place where the agent is supposed to act (or to abstain from acting) for instance, usually is connected with the act–description. It may be part of the subject–description, however. What’s the difference? Is “No person shall use obscene language in public places” equivalent to “No person in a public place shall use obscene language”? Do these wordings express one and the same norm? It seems so. Nevertheless, legal practice shows that they do not always have the same legal consequences. Their meaning therefore differs. A person who finds himself in a place that is not public (his house for instance) and uses an amplifier and a loud–speaker to send his obscene messages into the world, does not violate the norm expressed in the first formulation. But he does violate the second norm. (All this implies a theory of action according to which it is possible to act in a place that is not the place where one is).

This is just one aspect of the way in which the location of a norm–component within a norm may affect the meaning of that norm. Besides the question of the relation between the act–description and the indication of the place where the addressee is supposed to act, many similar questions arise. To what extent, for instance, is an indication of time to be considered as a norm–condition? The general problem is to gain some insight in the interaction between norm–components. What we are after, is some overall picture of the patterns that may be discerned in the wording of prescriptive norms, and of the consequences these patterns have in legal reasoning. Moreover, the theory should enable us to decide what patterns are equivalent, and are interchangeable without any loss of meaning.

The second main problem is the clarification of basic concepts. To illustrate the point and draw attention to the importance of semantical analysis, I will mention two problems. The first concerns the notion of ‘ought’, more specifically the question how the ‘ought-to-do’ is related to the ‘ought-to-be’. It is often assumed implicitly, among others by many deontic logicians, that a norm like ‘John ought to refrain from smoking’ is equivalent to ‘It ought to be the case that John does not smoke’. I doubt whether that is true. The latter norm, it seems to me, is not necessarily a prescription for John only. It may require others to do something to prevent John from smoking. In law, an ‘ought-to-be’–phrase may be chosen exactly for that purpose. In other cases, an ‘ought-to-be’–formulation is a clear sign of legislative errors. An example from Dutch law is the norm that prohibited a state of affairs that was described as: ‘more than three ships lie next to each other in the breadth of the river’. Since it was not clear who was supposed to act, and what exactly had to be done, it gave rise to some problems of statutory construction.

The second problem is whether ‘intentional’ can be considered as a predicate of actions. Is it possible, for instance, to recast ‘John’s action was an intentional killing’ as ‘John’s action was a killing and was intentional’? (Cf. [Prakken, 1993, p. 31]. Davidson [Davidson, 1982, p. 46] argues that this is impossible, because it gives rise to contradictions. Suppose that John killed a man who was the father of three children. By killing him, John therefore deprived these children of their father. Now according to Davidson there is only one action. The killing and the depriving are identical. Let us suppose further, that John intended to kill, but that he did not know that his victim was a father. In that case he did not intentionally deprive the children of their father. The conclusion that follows then, is that the killing was both intentional and not intentional.

Now whatever the solution to all of these problems may be, it is evident that the task of constructing theories of the basic legal knowledge categories (the abstract models as I called them) is far from easy. It requires, among other things, a theory of legal semantics, a theory of action and of act-descriptions and a thorough investigation of many examples from legal practice. In fact, some of this work has been done in legal theory. The results are promising. On the whole however, they do not suffice to supply those who work in the field of AI and Law with tailor-made solutions. Even legal theorists tend to have theories that are too remote from legal practice and their analyses often do not take into account the many nuances and complexities of legal knowledge categories. It is exactly the development of legal knowledge based systems that may provide an impetus to the development of more accurate theories.

## **6 Abstract models and logical form**

A considerable part of the work involved in constructing these theories may be accomplished without the use of logical tools. We are still in a phase in which consideration of the functions of legal norms and semantical inquiries into legal language may bring us further than we are now. If all this is lacking, formalization of law is a leap in the dark.

Nevertheless, one will reach a stage in the development of these models where a precise logical language may be profitable. After all, one may assume that many of the problems I mentioned earlier are problems of logical form. We may be able to detect these logical forms without recourse to the technical language of some logical system, as we often do. Even so, the use of logical tools may sharpen our awareness of them and give us the means to express the complexities that are hidden beneath the surface structure of law.

So it seems to me that if we need a precise account of the law that we would like to represent in legal knowledge based systems, the next step would be an attempt to capture its logical form.

Once again I would like to stress that this does not yet involve a choice for certain AI representation techniques at the expense of other techniques. The sole aim is to identify the object of representation. Of course, there is no identification without representation. One may, however, use logic primarily as an instrument to achieve an accurate description of the legal domain. Whether it is an adequate representation technique for legal knowledge based systems, as the logic programming approach claims [Kowalski & Sergot, 1990], is a different matter then, that remains to be seen. The idea behind this proposal is that we can only assess the accuracy of representation techniques in as far as we have a clear view of the domain. The quest for logical form may improve that view.

I have to explain what I understand by logical form, because that notion may be used in a broad, but also in a narrow sense. If we use the former, the logical form of a sentence is the form of that sentence after it has been translated into the language of a logical system. According to this view, logical form is relative in two ways. First, any logical form is related to a logical theory. Since there are many different logical theories, there are many different logical forms. The logical form of 'All cars are vehicles' may be 'p' if it is translated into the language of propositional logic. It may be 'for all x, x is a vehicle', if first order predicate logic is used (assuming that the range of application of

the variable  $x$  consists of the set of cars):

$$\forall_x(Vx)$$

Secondly, even within one and the same logical theory the logical form of a sentence may differ. It depends on the degree of refinement one aims at. The logical form ‘for all  $x$ , if  $x$  is a car, then  $x$  is a vehicle’ is more refined than ‘for all  $x$ ,  $x$  is a vehicle’:

$$\forall_x(Cx \rightarrow Vx)$$

The decision how much refinement is needed, is generally ad hoc. The purpose may be to highlight certain aspects of the sentence, and one may adjust the degree of refinement to that purpose. In this context, it may be noted that many of the representations of legal knowledge that we come across in the literature on legal expert systems are rather superficial. To a large extent the surface structure of legal language is simply reproduced. Article 53 of the Amsterdam APV, for instance, might have been represented as: for all  $x$ , if  $x$  is [Harassing other people in a public building] then  $x$  is [forbidden]:

$$\forall_x(Hx \rightarrow F^a x)$$

Now there is nothing wrong with simplification if it is one’s aim to stress certain aspects of legal knowledge, or certain aspects of the reasoning process. Deliberately simplified working models may be very useful. There is a problem, however, if we do not face the question what degree of refinement adequate legal knowledge representation requires.

Now there is another conception of logical form in which logical form no longer varies with the degree of refinement. The basic idea is that a sentence has an import – the set of sentences the given sentence is entailed by – and a purport – the set of sentences the given sentence entails. The notion of entailment is not yet related to a system of logic, but related to common-sense reasoning. The sentence ‘John kicked Jill’, for instance, is implied by ‘John kicked Jill with his left foot’, and by ‘John kicked Jill on her left knee which was already in bad shape’. On the other hand the sentence implies – among other things – : ‘somebody kicked Jill’, ‘Jill was kicked’, and ‘there occurred a kicking’. It does not imply that somebody was wounded or maltreated: both inferences depend on knowledge that is external to the sentence ‘John kicked Jill’. To give the logical form of a sentence then, is ‘to give its logical location in the totality of sentences, to describe it in a way that explicitly determines what sentences it entails and what sentences it is entailed by’ [Davidson, 1982, p. 140].

Now if we use first order predicate logic, as Davidson does, we will have to find a formula (a representation of the sentence ‘John kicked Jill’) that makes it possible to derive everything that can be inferred from ‘John kicked Jill’. The formula has to be ‘totally explicit’. The formula ‘Kicked(John,Jill)’, for instance, will not do. Predicate logic enables us to infer that somebody kicked Jill, but not that there was a kicking. In this respect the better formula is: ‘there is an  $x$ , such that  $x$  is a kicking, and  $x$  was done by John and  $x$  was done to Jill’, because it provides a ground for both inferences:

$$\exists_x(Kx \ \& \ Bxa \ \& \ Txb) \text{ with } a = \text{John } b = \text{Jill}.$$

There is no certainty, however, that this formula gives us *the* logical form of ‘John kicked Jill’. It is more like a hypothesis that has to be tested. I suspect this conception of logical form may be important for legal knowledge representation. It may function as its ‘regulative principle’ that provides ‘at once a standard of appraisal and a stimulus for action’ [Rescher, 1992, p. 118].

There is some doubt, however, whether we really need the complex representations that it brings. Consider article 53 of the Amsterdam APV. If we try to give the logical form of that regulation, the following formula is an approximation: 'for all w,x,y and z: if x is a person and y is a person and not (x = y) and z is harassing and z is by x and z is to y and z is in w and w is a public building, then z is forbidden':

$$\forall_w \forall_x \forall_y \forall_z [(Px \ \& \ Py \ \& \ \sim(x = y) \ \& \ Hz \ \& \ Bxz \ \& \ Tyz \ \& \ Izw \ \& \ Pw) \rightarrow F^a z]$$

Someone may say that this representation is too complex for the purposes of legal expert systems, and prefer a simpler version, like 'for all x, if x is harassing another person in a public building then x is forbidden' is:

$$\forall_x (Hx \rightarrow F^a x)$$

However, if knowledge representation is as easy as that, the most intricate rules of law may be represented in a very superficial manner. So the question is: How much of the logical form has to be unfolded?

## References

- T.J.M. Bench-Capon. Deontic Logic: Who needs it? In J.A. Breuker, editor, *Proceedings of the workshop Artificial Normative Reasoning at the European Conference on Artificial Intelligence*, pages 69–78, Amsterdam, 1994.
- P.W. Brouwer. *Coherence in Law: An Analytical Study*. PhD thesis, University of Leiden, Groningen, The Netherlands, June 1990. In dutch.
- D. Davidson. *Essays on Actions and Events*. Oxford, 1982.
- Ronald Dworkin. *Law's Empire*. Fontana, London, 1986.
- D.I. Gold and R.E. Susskind. Expert Systems in Law. A Jurisprudential and Formal Specification Approach. In A.A. Martino and F.S. Natali, editors, *Atti Preliminari del II Convegno Internazionale Logica, Informatica, Diritto*, Firenze, 1985.
- H.L.A. Hart. *The Concept of Law*. Second Edition, Clarendon Press, Oxford, 1994.
- O.W. Holmes. The Path of the Law. In R.A. Posner, editor, *The Essential Holmes*, Chicago and London 1992. From: 10 Harvard Law Review 457 (1897).
- Hans Kelsen. *Reine Rechtslehre*. Second edition, Franz Deuticke Verlag, Vienna, 1960.
- R.A. Kowalski and M.J. Sergot. The use of logical models in legal problem solving. *Ratio Juris*, 3(2):201–218, 1990.
- L.T. McCarty. Artificial Intelligence and Law: How to get there from here. *Ratio Juris*, 3(2):189–200, 1990.
- Henry Prakken. *Logical Tools for Modelling Legal Argument*. PhD thesis, Free University, Amsterdam, 1993.
- N. Rescher. *Ethical Idealism*. University of California Press, 1992.
- M. Smith. *Legal Expert Systems: Discussion of Theoretical Assumptions*. PhD thesis, University of Utrecht, The Netherlands, 1994. In dutch.
- G.H. von Wright. *Norm and Action*. Routledge and Kegan Paul, London, 1963.