

Legal knowledge based systems
JURIX '95
Telecommunication and AI & Law

The Foundation for Legal Knowledge Systems

Editors:

J.C. Hage

T.J.M. Bench-Capon

M.J. Cohen

H.J. van den Herik

Ronald Leenes and Jörgen Svensson , Large-scale computer networks and the future of legal knowledge-based systems, in: J.C. Hage, T.J.M. Bench-Capon, M.J. Cohen, H.J. van den Herik (eds.), *Legal knowledge based systems JURIX '95: Telecommunication and AI & Law*, Lelystad: Koninklijke Vermande, 1995, 75-86, ISBN 90 5458 252 9.

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



LARGE-SCALE COMPUTER NETWORKS AND THE FUTURE OF LEGAL KNOWLEDGE-BASED SYSTEMS

Ronald Leenes and Jörgen Svensson
*Department of Public Administration and Public Policy
University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands
email: {r.e.leenes, j.s.svensson}@bsk.utwente.nl*

Abstract

In this paper we investigate the relation between legal knowledge-based systems and large-scale computer networks such as the Internet. On the one hand, researchers of legal knowledge-based systems have claimed huge possibilities, but despite the efforts over the last twenty years, the number of practical applications actually in use is still small. On the other hand the computer networks, which in a practical form are only around for a few years, are gaining momentum and are expected to have a serious impact on many fields of society, including the legal domain. We look at the developments in both fields and ask the question: What can researchers of legal knowledge-based systems learn from the recent developments on the Net in order to improve the practical application of their research?

1 Computer support for the legal domain

Modern computers have lost the intuitive connection with the notion of 'computing'. With the help of computers and computer networks, we are nowadays able to acquire, store, retrieve, process, transport and present information which may be useful for all kinds of decision making. The value of computers as decision-support systems is now acknowledged in many domains, and many people in very different professions use decision-support systems. The possibilities of computers have also gained attention in the legal domain. Legal professionals are aware that computers cannot only serve as word processors, but may also support legal problem solving. With respect to this possibility of supporting legal professionals two very different developments can be noted.

On the one hand, for well over twenty years now, there is research on Legal Knowledge-Based Systems (LKBS). In this research, the idea of providing legal professionals with adequate information and knowledge is one of the driving forces. Researchers not only try to get more fundamental insight into legal reasoning, they also attempt to develop practical decision-support systems. The success of this applied research, in the form of the number of implemented applications is limited, to say the least. Only a few systems have found their way to the desk of the legal professional and even fewer have acquired a permanent position.

On the other hand, we see the rapid development of telecommunication via computers. The Internet (or the Net) is gaining more and more attention in society. And, although many may think the general attention for the Net is a hype, we all can see that the Net is a very promising development for the legal domain. Many legal professionals are already showing interest in the Net. And despite the fact that the Net in its current form only is around for about two years, it is already clear that it will be a matter of a few years before the use of the Net by legal professionals will be standard practice. This makes us, researchers in LKBS, think. Why does the Net gain momentum and are LKBS staying behind? And what can we do to make our research of more practical use?

In this paper we try to answer these questions. In order to do so, we start the next section discussing the developments in both technologies. Then, in section 3, we compare the technologies on several relevant dimensions. Based on this comparison, in

section 4, we come to our recommendations for applied LKBS research. And finally, in section 5, we give our conclusions and make some additional observations.

2 LKBS and the Net: an overview

Legal knowledge-based systems and the Net are the two relatively new technologies. We discuss them both and see how they have developed.

2.1 Legal Knowledge-Based Systems

Artificial Intelligence (AI) focuses on capturing the idea of intelligent reasoning in computers and on making “intelligent” computer programs. In the early days of AI research, it was recognised that intelligent performance of both humans and computers requires a great deal of factual knowledge. This means that the development of systems which demonstrate a general artificial intelligence is practically impossible, because it presupposes the development of enormous knowledge bases. It was thus recognised that AI research should aim at developing more specific systems, confined to knowledge domains with limited zones. The limitation of a domain reduces the amount of knowledge to be incorporated in the system and makes it possible to aim for systems which can make expert-like decisions in that domain.

‘Expert systems’ thus are computer programs which pertain to knowledge domains of limited size. The early results in expert systems such as MYCIN (Buchanan and Shortliffe, 1984) and PROSPECTOR (Hayes-Roth, 1987, p. 287) proved that this idea was very promising. It was believed that expert systems had a glorious future and that they would enter many domains.

The early legal expert systems

The advances in AI were also expected to produce results in the field of law (Buchanan and Headrick, 1970). It was thought that the formal nature of many legal domains was very suitable for representation in expert systems. High hopes were expressed about the possibilities of computer programs which could perform tasks only associated to legal experts until then. Research focused on capturing legal reasoning in computer systems and on finding suitable ways of representing legal knowledge.

This research into legal expert systems led to several prototype systems of which the most well-known examples are: The rule-based model of the British Nationality Act (Sergot, 1988), TAXMAN (McCarty, 1980) and LEGOL (Stamper, 1980). These systems were in the first place research systems, however, and not yet full-grown commercial applications. As an offspring of these experiments, a next generation of legal expert systems was developed. Well-known projects of this second generation are: The Alvey Demonstrator Project (Bench-Capon, 1991), JURICAS (De Mulder, 1984; van Noortwijk and Stubbe, 1986), PROLEX (van den Berg *et al.*, 1987) and TESSEC (Nieuwenhuis, 1989).

Although these new systems also started off as research projects, they had the pretension to be practical legal expert systems. Unfortunately, they failed in this respect. Although all the systems showed the practical possibilities AI and Law had to offer, and were received with some enthusiasm by the public, the research community was eager to condemn them. It was said that these systems, of which the developers were so proud, were “not really intelligent” and could not solve the problems legal professionals were really interested in, like the so-called hard cases.¹ The systems also lacked fundamental aspects of real legal knowledge and real legal reasoning like open texture and defeasible reasoning.

¹ See Hage *et al.* (1994) for an account about the distinction between hard and complex cases.

Large-scale computer networks and the future of legal knowledge-based systems

Knowledge-based systems' research: a diversification

After the phase of these first legal expert systems, the research into AI and Law diversified. A clear indicator for this diversification is the new name given to the field: 'Legal Knowledge-based Systems'. The developments are best noticed when browsing through the proceedings of the International Conference on Artificial Intelligence and Law of 1991 (38 papers) and 1995 (35 papers).²

First, we see in this period a decline in papers about practical systems (25 in 1991, 14 in 1995). Second, the fundamental research into logics underlying legal reasoning is increasing (5 papers in 1991, 11 in 1995). Topics included in this type of research are: deontic logics, non-monotonic logics, dialectical logics etc. Third, we see a growing interest in more advanced schemes for representation and retrieval of legal knowledge. Papers in conceptual representations developed from 4 papers in 1991 to 7 papers in 1995. An interesting instance of the early research is (Dick, 1991). Fourth and finally, we see a growing attention for the automated analyses of legal sources. Under this type of research we may classify statistical analysis of legal text material (Sweighofer *et al.*, 1995) and the developments in neural networks (Walker *et al.*, 1991).

Much of the current research is no longer directly aimed at the development of practical systems, but at the gathering of theoretical insight into the possibility of modelling legal knowledge and artificial reasoning. The AI-and-Law community has recognised that making real intelligent systems is very difficult and (possibly) not rewarding enough.

2.2 Telecommunication with computers: getting information over the Web

In the last years we have seen a rapid development of computer networks. Personal computers in the office are being integrated in local area networks, which are connected to wide area networks, like the Internet. The connection of computers to these types of networks offers new possibilities for communicating information. These possibilities are considered so important that people now talk of the Information Superhighway. We will look at the general development of communication via the Net and at the meaning of this development for the legal professional.

The development of 'The Net'

Until about 1993, the Internet consisted only of a large number of interconnected computers, mainly at universities and research institutes. Although there was much information available on the Net, it required some Zen of Internet to find it. The programs with which Internet could be accessed were difficult to understand and access to information sources on the Net required at least a basic understanding of the operation of Mainframes and Minis.

The development of the World Wide Web (WWW, or the Web) marks a change in this situation. The Web basically offers a way to make documents available on the Net in a standardised form, so that they can be accessed with the help of a standard 'Web browser' (e.g., MOSAIC and NETSCAPE) which are available for all current types of operating systems and which are easy to use.

Any person who is connected to the Net and who wants to offer information to other users on the Net now can set up a Web Server and become an 'Information provider'. All one has to do is use the simple HTML (HyperText Markup Language) to format one's documents. HTML basically offers two possibilities to the provider:

- a. HTML can be used to make documents look good (for instance, it offers the capability to include multimedia features),

² Inventory made by the second author.

- b. HTML makes it possible to link documents in a hypertext manner by means of 'Hyperlinks'.

Moreover, HTML gives the real power to the Web. Hyperlinks allow the provider to define 'hot-spots' in a document which in turn allow the user to jump from one document (which may be on any computer on the Net) to another document (which may be on any other computer on the Net), just by clicking on that 'hot-spot' with the mouse. In practice this means that anyone with a modem or a direct connection to the Net can enter the Web and jump from document to document when searching for interesting information.

The growth in number of information providers is now enormous. The Web has developed from almost nothing to a huge system containing millions of documents in only two years.³

In order to help the user to find information on the Web, so called 'search engines' have been developed which index the Web by crawling or travelling through the documents and which provide the possibility to search indexed documents on keywords. Another service for searching users is offered by the numerous list-of-lists available on the Web.

One of the major reasons why commercial publishers still refrain from publishing information of commercial value on the Web is a financial one. Publishers want to be paid for their services and that is something not settled yet. However, a number of experiments to solve the payment problem are now under way, and once this issue is settled publishers certainly will enter the Web.

The Net as an information source for the legal professional

It is clear that the Net (in its prominent form of the Web) is an interesting development for the legal professional, who is essentially a knowledge worker. In order to make adequate decisions in the legal domain, knowledge of that domain is of the utmost importance. This means on the one hand that a legal professional has to have a general knowledge of the legal field (which can be acquired through vocational training and training on the job). On the other hand, however, it means that (s)he has to have specific and up-to-date information with respect to the legal questions at hand, in particular concerning statute law in force in a certain domain and the developments in case law.

Until now, this information had to be collected in large local libraries and through subscription to juridical handbooks.⁴ The Web offers the possibilities to search for legal information all over the world and thus reduces the need for maintaining these libraries and hand books. Although at this moment the amount of legal information on the Net is still rather limited,⁵ both government agencies and publishers are expressing their interest to start publishing legal information on it. In only a few years the Net will play a major role in communicating legal information.

3 The Net and LKBS compared

LKBS and the Net have in common that they compete for the attention of the legal professional. From the descriptions above, it is clear that there are huge differences in what they offer and in the way they compete for attention. In order to find the factors

³ According to the Lycos FAQ (June 15, 1995; <http://www.lycos.com/lycos-faq.html#q1a>), there are about 5 million documents available on the Web, with an average text file size of almost 8.000 characters.

⁴ Electronic legal databases such as WESTLAW and LEXIS in the US and NLEX, KLUWER JURIDISCHE DATABANK en PARAC in The Netherlands only have played a minor role in reducing this need for paper based documents.

⁵ The German University of Saarbrücken, and the American Cornell Law School for instance both maintain pointers to relevant legal material on the Internet.

which influence the attention given to both the Net and LKBS in the legal domain, we will now make some comparisons with respect to several dimensions. These dimensions are:

- the type of support to the user,
- the overall benefit for the user (intended and realised),
- the practice of system development,
- the relationship (appeal) to the user.

We summarise our comparison in Table 2.

3.1 *Type of support to the user*

Comparing LKBS and the Net in the context of their contributions as decision-support systems for the legal professional, the first distinction is the type of support. We identify four different types of support:

1. Computers and networks may play a role in the *communication* between people. Especially the Net aims at the exchange of information without the need for simultaneous presence of sender and receiver. Information can be exchanged over the Net in various formats which can be processed by the recipient, in contrast to for example fax messages.
2. Both the Net and LKBS are providing means for *storing* information and knowledge. On the Net this is generally done in the form of documents and databases. In LKBS this may also be done in the form of different knowledge representation formalisms. Both types of technology also provide possibilities to *retrieve* stored information, where a special quality of the Net is the ability to search for knowledge which is distributed over many systems, while LKBS research is more directed at intelligent search and matching.
3. *Structuring and formalising* available information can support human decision making. Computer tools can be of assistance in this task. Especially in LKBS this possibility is the topic of research. Examples along this line are the norm frames as developed in (Van Kralingen *et al.*, 1993), Reason-Based Logic (Leenes *et al.*, 1995) and the findings in the EXPERTISZE project (Kordelaar, 1993).
4. LKBS research and especially the research into legal expert systems has focused on *providing reasoned advice* about problems at hand. A so-called expert system can use a case description as input and produce a 'solution' to the case based on the case data and the knowledge represented in the system.

Table 1 shows the different types of support and the suitability of both LKBS and the Net to deliver this support. The Net is particularly suitable for communication support and for storing and retrieving information. LKBS research has always focused on automated reasoning, although it did not succeed too well in this aim (hence only one '+'), but it also is considered to have some virtues for the other three types of support.

	Communicate	Store and retrieve	Organise	Reason
The Net	++	++	-	-
LKBS	+	+	+	+

Table 1: Focus of LKBS and the Net on types of support

3.2 The overall benefit for the user (intended and realised)

The intended benefit for the user of both types of systems is rather different. On the one hand, LKBS research has especially aimed at developing reasoning systems and it has done so in an ambitious manner. The notions of 'Artificial Intelligence' and 'Expert Systems' convey the idea of developing systems which can reason as human experts, or even better. However, the idea that LKBS can take the place of lawyers and judges has proved far too optimistic. The legal expert systems available today cannot keep up with the promise. The systems generally concern limited and very well-structured domains in which they can deal with help to solve standard cases. When the problem at hand crosses the boundaries of the domain, or when the case proves to be a 'hard case', LKBS do not perform well, and the outcome of the reasoning process soon becomes worthless. Apart from this practical problem, the strain between the ambitious goals and the actual practice of LKBS produces serious problems. As discussed, LKBS may perform adequately in limited and well-structured domains (as we can find for instance in many tax and social insurance schemes). However, the failure of LKBS to perform well in general, seems to hinder the systems from entering the market. Because AI still is the holy grail, even if a system is objectively of practical use, it may be refused on the grounds that it does not meet the high standards of LKBS. (JURICAS is not a real expert system but just a decision tree, TESSEC has a naive knowledge representation and Neural Nets cannot explain their reasoning).

The Net on the other hand, raises only expectations about communication and available information. It offers information, which the user has to interpret and apply herself. Although the amount of useful legal information at the moment is still limited, we see improvements every day. Once information about cases and statutory provisions becomes available, the Net will be the information source for legal professionals. The danger that raised expectations are not met is rather small. If the user does not find the information she is looking for, she has to look somewhere else, or she has to settle for something else. Browsing the Web in this respect resembles searching for a book in a library. You know you may not always find what you are looking for, but as long as there is a fair amount of useful information you will come back.

3.3 Development and maintenance

Large differences can also be seen in the ways LKBS and the Net are being developed and maintained. These differences are related to one major difference between the technologies, namely in the level of coding needed to incorporate knowledge and information in the system.

A statement that holds for both LKBS and the Net is that information and knowledge to be incorporated in the system, must be available in a form understandable to the system and compatible with the type of operations the system is supposed to perform. This means that the developer of a legal information system has to translate legal knowledge from its original form (the form in which it is found, e.g., written texts) into the format required by the system.

Especially in the development of rule-based systems (still the major technique for developing practical systems) the aim of automated reasoning poses important problems in this respect. In order to enable automatic reasoning with knowledge it is necessary to interpret and formalise the knowledge in sophisticated manners. Rule-based expert systems therefore typically require a labour-intensive coding of knowledge which has to be performed by a coding specialist or 'knowledge engineer'.⁶ In addition to this

⁶ The idea of using expert system shells to speed up the process of system development, although very useful, does not help much to solve this problem. With the help of a shell, the process of developing a working system is still quite hard, and in practice it is found that, due to the fact that every developer has his particular wishes, there are almost as many shells as there are knowledge-based systems.

Large-scale computer networks and the future of legal knowledge-based systems

problem, the knowledge base of a rule-based system needs to be consistent; the different pieces of knowledge have to be coded in relation to each other. This need for consistency means in practice that it is very difficult to work on a system with more than a few people because every member of the developing team has to be in close contact with the others in order to ensure that his piece of coded knowledge fits in with the rest. It should however be acknowledged that other LKBS techniques like neural nets and case-based reasoning systems are less affected by the problem of coding (in both instances only a much simpler level of manual coding suffice) and also the research into automated analyses of legal texts holds a promise here.

As we stated before, the Net has an obvious aim. The system does not have to reason with the knowledge; it only has to be able to contain knowledge and pass it through to the user. Because of this more limited aim, knowledge and information on the Net can be more loosely related, and relatively simple coding schemes suffice. For the HTML tagging scheme, this coding can even be done automatically by computers. And, because much legal information nowadays is available in electronic form, it is easy to transform this information into Web pages. The simple coding requirements and the low consistency requirements provide the Net with four major advantages:

- The Net can develop incrementally (new knowledge can be added easily);
- almost everybody can add knowledge to the Net;
- adding knowledge to the Net is relatively inexpensive;
- the process of adding knowledge to the Net is fast, meaning that the system can be held up-to-date.

For instance, in the Hermes project, rulings of the supreme court are made available on the Internet for free, almost immediately after they were passed.⁷

3.4 The additional incentives to adopt the technology

The last difference between LKBS and the Net is partly connected to the differences described above. It concerns the additional incentives for the user and the developer to adopt the technology. These additional incentives differ on two dimensions: investment and sexiness.

In the first place, investment in LKBS is still a matter of buying single applications. This means that a legal professional who deals with several legal specialities has to invest in several separate systems. The Net on the other hand, is a one time investment only. If the legal professional decides that she needs the Net, she will get the whole technology and will not need to think about this again (she only has to register separately for commercial services).

Secondly, the Net has a bit more sex-appeal. In LKBS the interaction with the user is mainly text-based: the user types in information and sees the responses of the system on his monitor. The Web, however, combines pictures, sounds, movies, text and thus can be used to present information in an attractive form. Moreover, the user may discover the world by travelling the virtual world which the Net actually is. Although one might say that in the legal domain content comes first, we believe that sexiness can certainly influence a decision to invest in technology.

	keyword	The Net	LKBS
--	----------------	----------------	-------------

⁷ Project Hermes, see <http://www.law.cornell.edu> for a starting point.

type of support		- communicate (++) - store and retrieve (++)	- communicate (+) - store and retrieve (+) - organise (+) - reason (+)
benefit for the user	level of support	providing information about all sorts of legal questions	intelligent assistance in simple cases in a strictly limited domain
development	coding	simple, fast	difficult, slow
	development team	multiple providers supplying one system	individuals or small groups making separate systems
incentive to adopt	investment	one time only decision	separate investment decisions
	sexiness	- multimedia - global communication	- text and numbers - local PC

Table 2: Characteristics of Web and LKBS compared

4 Making LKBS more practical

In the previous section we have compared LKBS and the Net with respect to their practical meaning for legal decision support. We used four dimensions to compare both developments: the type of support to the user, the overall benefit for the user, the practice of system development and the incentive to adopt the technology. This comparison has led to several insights, based on which we now can formulate five recommendations for the LKBS community.

First, we have seen that LKBS research concerns a whole range of support types. It is not limited to the development of expert systems. It also addresses the problems of storing and retrieving legal knowledge, of organising legal knowledge and of communicating about legal knowledge. However, where the development of practical expert systems in the past has received much attention, the other types of applications only have been researched in small laboratory situations. In our opinion, the time is now ripe to develop these ideas into practical systems.

Second, in our view, the legal expert systems of the first generations were not really failures in a technical sense but they were made failures both by the hyped expectations (really intelligent systems) and the unreal ambitions of the LKBS research community. Both researchers and potential users of LKBS products should be more aware of this and should learn that really intelligent systems will never exist, at least not in our lifetime, and that practical applications for limited and well-structured domains are already there. Although the AI component in these systems may be limited, they are spin-offs of the LKBS research which proves that there really is something to LKBS. Furthermore, the use of such systems can provide scientific knowledge about the way decision support can be given by LKBS.

Third, although fundamental research in AI and Law should still focus on developing coding techniques which really allow intelligent legal reasoning, applied LKBS research should aim at developing coding schemes which are practical. They should find an equilibrium between the intricate coding techniques used in LKBS of today and the perhaps too simple and too general technique of HTML tagging. New coding techniques should reduce the necessity of special coding experts (because there will always be too few of them) and allow more people to work together on a single project. Only if we find

such coding techniques, it will be possible to improve systems which can be kept up-to-date.

Fourth, LKBS researchers can profit from the successful aspects of the Net to promote their products in a direct way: LKBS can join the Net. This can be done in several ways.

One first step may be to use the Net as a carrier for LKBS. Not only is it possible to spread LKBS products via the Net, it is also possible to construct systems which can be accessed over the Net. By using a Web page as an interface, LKBS can be made available to users all over the world who have a Web browser.

One further step may be to integrate LKBS and the Net. As Soper and Bench-Capon (1992) have pointed out, expert-system technology can be enhanced with hypertext and hypermedia techniques (techniques which now form the basis of HTML and the popular WWW browsers). Lenk (1990) foresees even more integrated systems, in which a computer can service the user with normal information retrieval (as on the Web), LKBS and on-line connections with legal professionals in some distant back-office.

Finally, the Net is becoming an important source of legal knowledge. LKBS researchers should research the possibilities of applying this source in their system development. Especially the possibility of (semi-) automated acquisition of legal knowledge for knowledge-based systems is relevant in this respect. As more information becomes available on the Net, it becomes possible to analyse this information with statistical methods. The SALOMON project at CU Leuven goes some way along this line (Moens *et al.*, 1995). It may well be that the real future of AI and Law lies in this direction. The results of automatic knowledge extraction from large amounts of data are promising. Although the construction of legal knowledge-based systems capable of performing 'real' legal reasoning is yet far beyond our capabilities, we know that one of the main obstacles for this goal is the amount of knowledge needed to do the problem solving. Until now knowledge representation of large knowledge bases was done by hand. When large amounts of legal information become available on the Net and we succeed in developing (semi-) automated knowledge acquisition, who knows where we get.

5 Conclusion

For us, researchers in LKBS, the Net sometimes seems something to fear. The Net shows a very impressive growth and it seems that it has delivered more in the last two years than research in AI and Law has in the last twenty years. The Net and its development draw increasing attention, and at the same time we see a declining interest in LKBS. Our view is that we should not fear the Net, but investigate its success and try to learn from it. In this paper we have done this by comparing both technologies on several dimensions.

We found that LKBS research aims at a broader and more ambitious type of support for legal professionals than the Net, and that it is exactly this ambition that has made it almost impossible to achieve success. The expressed ambition leads to unrealistic expectations from intended users, and the researchers themselves are never satisfied with their products. Moreover, the ambition of intelligent automated legal reasoning introduces a need for intricate coding schemes, which makes that legal expert systems are limited to relatively very small and structured knowledge domains.

From our observations we arrived at the following recommendations for developers of LKBS:

1. Look at *all* the benefits that LKBS can offer. LKBS include not only traditional expert systems but also systems for communicating about knowledge, storing and retrieving knowledge, and organising knowledge.
2. Get rid of the AI aura. Although this aura may have done the research some good twenty years ago, it is now seriously hindering practical development.

3. Develop simple but sufficient coding schemes for rule-based systems which will allow more people to work on them.
4. Use the Net as a carrier for your products. Distribute your system over the Net, develop Web interfaces and try to integrate your LKBS with other services on the Net.
5. Try to develop techniques to make the best use of the enormous amount of legal information which is now becoming available on the Net.

Finally, although this paper may have conveyed a different message, LKBS researchers should concentrate on the specific added value of LKBS in field of automated legal reasoning. The Net is now leading to a real information explosion, but it gives the user few tools to handle the information overload it presents. As Simon (1976) and Silver (1991) argue, an increase in the amount of information available to a decision maker does not eliminate the information problem. It rather leads to a shift in needs. Instead of more information, people will want systems which support their selection of information and systems which really support decision making. LKBS research may provide these systems. And so, as the Net increases the amount of legal information available, it also increases the demand for real LKBS!

References

- Bench-Capon, T.J.M. (1991). *Knowledge Based Systems and Legal Applications*. Academic Press, London.
- Berg, P.H. van den, Opdorp, G.J. van, Oskamp, A., Prakken, H., Schrickx, J.A., and Walker, R.F. (1990). *PROLEXS, een ontwikkelomgeving voor juridische expertsystemen* [PROLEXS, a toolkit for developing legal expert systems], final report SPIN-PKO. Computer/Law Institute, Vrije Universiteit, Amsterdam.
- Buchanan, B.C., and Headrick, T.E. (1970). Some Speculations About Artificial Intelligence and Legal Reasoning. *Stanford Law Review*, 1970 (40).
- Buchanan, B.C., and Shortliffe, E.H. (eds.) (1984). *Rule-based expert systems : the MYCIN experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading (Mass.).
- Dayal, S., Harner, M., Johnson, P., and Mead, D. (1993). Beyond Knowledge Representation: Commercial Uses for Legal Knowledge Based Systems. *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pp. 167-174. Association for Computing Machinery, New York.
- Dick, J.P. (1991). Representation of legal text for conceptual retrieval. *Proceedings of the third International Conference on Artificial Intelligence and Law*, pp. 244-253. Association for Computing Machinery, New York.
- Hage, J.C. (1987). Themis als robot. *RM Themis*, 1987 (5), pp. 238-248.
- Hage, J.C., Leenes, R.E., and Lodder, A.R. (1994). Hard Cases: A Procedural Approach. *Artificial Intelligence and Law*, Vol. 2, pp. 113-167.
- Hayes-Roth, F. (1987). Expert Systems. *Encyclopedia of Artificial Intelligence* (eds. S.C. Shapiro and D. Eckroth). J. Wiley and Son, New York.
- Herik, H.J. van den (1991). *Kunnen computers rechtspreken?* Gouda Quint, Arnhem.
- Kralingen, R. van, Oskamp, O., and Reurings, E. (1993). Norm frames in the representation of laws. *JURIX '93 Legal Knowledge Based Systems: Intelligent Tools for Drafting Legislation, Computer - Supported Comparison of Law*, pp. 11-22. Vermande, Lelystad.
- Leenes, R., Lodder, A., and Hage, J. (1995). A dialog game for legal arguments. *Law, Computers and Artificial Intelligence*.
- Lenk, K (ed.) (1990). *Neue Informationsdienste im Verhaeltnis von Buerger und Verwaltung*. Decker & Mueller, Heidelberg.

Large-scale computer networks and the future of legal knowledge-based systems

- Mathijssen, L.J. (1995). An Intelligent Interface for Legal Databases. *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pp. 71-80. Association for Computing Machinery, New York.
- McCarty, L.T. (1980). The TAXMAN Project: Towards a Cognitive Theory of Legal Argument. *Computer Science and Law* (ed. B. Niblett), pp. 23-43.
- McCarty, L.T. (1977). Reflections on TAXMAN: An Experiment in Artificial Intelligence and Legal Reasoning. *Harvard Law Review*.
- Moens, M.F., and Uyttendaele (1995). Salomon: automatische extractie van correctieelonnissen (report of a discussion in the JURIX meetings). *NVKI-nieuwsbrief*, Vol. 12, No. 3, pp. 70-72.
- Mulder, R.V. de (1984). *Een model voor Juridische Informatica*. Vermande, Lelystad.
- Nieuwenhuis, M.A. (1989). *TESSEC: een expertsysteem voor de Algemene Bijstandswet*, doctoral thesis. Universiteit Twente, Enschede.
- Noortwijk, C. van, and Stubbe, A.L.W.M. (1986). *JURICAS: Erfrecht, een juridisch computeradviesstelsel*. Vermande, Lelystad.
- Opdorp, G.J. van, Walker, R.F., Schrickx, J.A., Groendijk, C., and Berg, P.H. van den (1991). Networks at Work: A Connectionist Approach to Non-Deductive Legal Reasoning. *Proceedings of the Third International Conference on Artificial Intelligence and Law*, pp. 278-287. Association for Computing Machinery, Oxford.
- Rissland, E.L., and Daniels, J.J. (1995). A Hybrid CBR-IR Approach to Legal Information Retrieval. *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pp. 52-61. Association for Computing Machinery, New York.
- Sergot, M. (1988). Representing Legislation as Logic Programs. *Machine Intelligence* (J.E. Hayes, D. Michie and J. Richards), Vol. 11, pp. 209-260.
- Silver (1991). *Systems that support decision makers: description and analysis*. J. Wiley & Sons, Chichester.
- Simon, H.A. (1976). *Administrative Behaviour* (3rd edition). The Free Press, New York.
- Soper, P., and Bench-Capon, T. (1994). Coupling Hypertext and knowledge based systems: two applications in the legal domain. *Artificial Intelligence and Law*, Vol. 2, pp. 293-314.
- Stamper, R. (1980). Modelling Legal Rules By Computer. *Computer Science and Law* (ed. B. Niblett), pp. 45-71.
- Svensson, J.S., Kordelaar, P.J.M., Wassink, J.G.J., and Eind, G.J. van het (1992). EXPERTISZE. A Tool for Determining the Effects of Social Security Legislation. *Legal Knowledge Based Systems: Information Technology and Law, JURIX '92* (eds. C.A.F.M. Grütters, J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and C.N.J. de Vey Mestdagh).

