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

D. Hunter, *Looking for Law in all the Wrong Places: Legal Theory and Legal Neural Networks,* in: A. Soeteman (eds.), Legal knowledge based systems JURIX 94: The Foundation for Legal Knowledge Systems, Lelystad: Koninklijke Vermande, 1994, pp. 55-64, 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

Looking for Law in all the Wrong Places: Legal Theory and Legal Neural Networks

Dan Hunter
Law School, University of Melbourne
Parkville Victoria 3052
Australia
eMail: dah@rumpole.law.unimelb.edu.au

Abstract

In this paper the author argues that neural networks have been poorly served by their applications to date in the legal domain. He argues that law and legal theory have been equally ignored in these applications. This paper thus raises a number of specific legal theoretical concerns. These include the difficulties that the inherent nature of neural networks present, the problems of using statistical, sub-symbolic techniques in a (purportedly) symbolic field, the dangers in choosing the training set, and the legal philosophy apparent in current neural network implementations. The author suggests how we might create neural networks in the future which, perhaps, will not give rise to such concerns.

1 Introduction

Fashions change, in artificial intelligence as in law. At one time neural networks were thought to contain the promise of truly intelligent machines. Their perceived similarity to human cognition marked them as different, new and full of hope. When it was found, however, that these networks could not perform the simplest of symbolic logical proofs they were consigned to the sad box of tricks at the back of the artificial intelligence researcher's cupboard: the dusty one marked 'Paradigms Lost'.

Time has caught up with neural networks, and now with new and more powerful learning algorithms, neural networks form a major and vibrant artificial intelligence sub-discipline. This has not gone unnoticed in law and neural networks have been built in legal domains for at least the last five years.

What are we to make of these legal neural networks? The modest aim of this most modest paper is to look at the question from a legal theoretical perspective. The paper asks to what extent the currently implemented legal neural networks have been built with any appreciation of legal theory. It makes a series of comments about the legal theoretical pitfalls into which some researchers have fallen. In the conclusion it suggests how we might avoid these pitfalls in future, including a proposal for a number of experiments in legal neural networks.

It is not the place here to review neural network basics. Readers unfamiliar with neural networks are referred to standard references such as [Andersen & Rosenfeld; Luger & Stubblefield, 1993; McClelland *et al.*, 1988; Michalski *et al.*, 1983; Mital & Johnson, 1992; Winston, 1992; Zurada, 1992].

2 Legal theoretical views of neural networks

A number of attempts have been made to apply neural network technology to legal problems. The major ones reported are [Rose & Belew, 1989; Rose & Belew, 1991; Warner, 1990; Warner, 1992; Warner, 1993; Bochereau *et al.*, 1991; Philipps, 1989; Philipps, 1991; Opdorp & Walker, 1990; Opdorp *et al.*, 1991; Walker *et al.*, 1991; Birmingham, 1992; Bench-Capon, 1993; Hobson & Slee, 1994]. Each of these implementations raise a number of legal theoretical issues. However, there is not space here to describe these systems. Instead, I look generally at the concerns which these systems have presented.

There are a number of legal theoretical issues presented by current neural network implementations. These fall into two basic areas – the nature of neural networks and the nature of the training set. We will see that the many neural networks approaches are not well suited to the legal domain. This is not to say that neural networks can never be appropriate for law, but rather that they are inappropriate for what the purposes of their implementors. Secondly, we will see that researchers must be extremely careful in the choice of the training sets used in legal neural networks. Since in neural networks we do not have the luxury of examining a symbolic representation of the law, we must be diligent in assessing the validity of the only information we have – the training set. An inappropriate choice of data will result in useless results.

2.1 The nature of neural networks

Neural networks have performed well in classification and recognition of 'objective' sensory data. The most successful field has been visual perception, where pattern recognition works extremely well and also it seems the basis for animal vision, [Winston, 1992]. The statistical basis of the paradigm means that it is very good at making correlations between a new pattern and a previously trained one. In law we could expect neural networks to perform similar tasks equally successfully, and we are not disappointed. Neural networks can perform useful pattern matching by 'recalling' a previous identical case to the one at bar, [Hobson & Slee, 1994]. Further, the network's statistical basis allows it to recognise associations between related cases by increasing the weighting on the links which correspond to these cases. In this way the neural network can perform classification-type processing. If presented with a sufficient number of cases which contain similar attributes and values, together with relevant outcomes, the neural network can classify these cases as being of one type. For example, 1000 cases containing the inputs of `offer` and `acceptance` with outcome of `valid_contract` would create a strong classification regime. Then, when presented with a new case (`offer = yes` and `acceptance = no`) it will 'recognise' this case as falling within this classification and return the given outcome (`valid_contract = no`).

Is legal reasoning simply a process of pattern recognition or classification? Unfortunately it is not, any more than it is a simple process of symbolic deduction. If legal reasoning were classification, then each case would fall within a particular classification or classifications, and a given result would follow. Since any practising lawyer can present alternative arguments on either side we know this is not all there is to legal reasoning. We do not have to choose between [Hart, 1961]'s reduction of hard cases to core and penumbra, or the critical legal theorists' view of indeterminacy [Kennedy, 1986; Kennedy, 1991], or other more radical postmodern views [Fish, 1982; Fish, 1983; Mackinnon, 1989]. We simply need to recognise that in law generating only one answer is insufficient, whether that answer is generated by logical deduction or by pattern matching. The paper of Hobson and Slee and those by Warner fail to recognise that classification is not reasoning.

The problem with classification is not the only difficulty in claims that neural networks can perform legal reasoning. Related to the preceding discussion, there is to be an assumption by some researchers that these systems can perform analogical reasoning. By analogical reasoning we mean the ability to retrieve precedents, adapt them better to fit the current case, and draw conclusions from the relationship [Burton, 1985; Levi, 1949; Ashley, 1990; Kolodner, 1993]. Can neural networks perform analogical reasoning? [Hobson & Slee, 1994] have, as a premise of their discussion, that their neural network system is to be used for case based reasoning; that is, for analogy. [Philipps, 1991, p. 998] shows a diagram that indicates a number of 'conclusions drawn by analogy' when these conclusions appear to be simply based upon classification.

To say that a classification or generalisation process produces results based on analogy is misleading. It falsely suggests, at least to this author, that the neural network can modify its conclusion by a process of case adaptation, when in fact no implementation has yet attempted this. Whether machine learning systems can be used in this way is a question which we are currently studying.[1] But it is quite clear that this is not what has been undertaken to date, and for [Hobson & Slee, 1994] and [Philipps, 1991] to claim so is misleading. Other researchers who use neural networks for classification do not present such problems.

If these were the only concerns about the use of neural networks then we could be sanguine about their future development in legal domains. However other broader claims are made. One author suggest that the nature of neural networks is such that it is not only useful for legal reasoning, but more grandly that it somehow represents the entirety of law. [Warner, 1993, p. 136], a law professor, is perhaps the boldest in his view of law:

> "While our language dictates a sequential description of the (legal reasoning) process, the process is in fact parallel. Many aspects of the problem resolution process are carried out simultaneously. The problem domain is defined by the initial statement of the problem. That initial problem is then resolved into a number of issues...the solution to which will be sought within the problem domain utilizing a subsymbolic paradigm that is not rule based" (references omitted)

A grand theory of parallelism indeed, singularly lacking in any authority, examples or other backing for the argument presented. Other authors are somewhat more

[1] This research is being undertaken at the Collaborative Law and Artificial Intelligence Research (CLAIR) project, at the University of Melbourne Law School and the Department of Computer Science and Computer Engineering, La Trobe University. The project examines rule induction systems as well as neural networks. It will be the subject of subsequent reports.

circumspect, though still advance interesting claims for neural networks. For example, [Rose & Belew, 1989, p. 140] and [Rose & Belew, 1991, pp. 2–3] suggest that the legal system contains elements of both localised and global concepts. That is, legal doctrine stems from the interaction of the corpus of judge–made and statute law. As [Rose & Belew, 1991, p. 2] put it, "Thus global concepts emerge from the interaction of a large number of localised decisions." They then suggest that this bears a striking similarity with the symbolic and connectionist approaches of artificial intelligence research. Hence, they argue, an integrated connectionist and symbolic architecture is appropriate for their legal model and retrieval system.

[Rose & Belew, 1991, p. 3] argue further that the ability of SCALIR to adjust its weights by back propagation accords with legal realists' view that "...law can(not) be adequately explained by some set of rules or concepts." I do not take issue with their artificial intelligence formalism, but I am concerned about their adoption of such a simplistic view of legal realism. Rose and Belew are in fact conflating the myriad strands of all descriptive jurisprudence with one type of statistical analysis. It may be true that statistics has been, and will continue to be, used by some legal realists, but that is not to say that all realist jurisprudence is limited to statistics. They are in effect viewing law as a statistical process, rather than a process of individual cases, interlocking obligations, relationships of power, race, and so on. It seems, at least to me, too simple a model to be accepted without further study. And certainly a view of legal realism of which many legal realists would be unhappy.

Rose and Belew might be accurate in arguing that the use of neural networks in *retrieval* accords with a legal realistic view of law. This seems reasonable, given that their SCALIR system is designed only for document and information retrieval. However, this obscures the vital point that neural networks do not accord with a legal realistic view when they are used for *reasoning*. Neural network conclusions are based on a statistical analysis of a training set. Assuming the training set is valid – a question addressed later – the output generated is simply a conclusion based upon a statistical weighting of the importance of the inputs. I would argue that one cannot suggest that all law is necessarily based upon statistics. Appeals to statistical data have often been rejected by courts, and represent at best one way of analysing a given legal domain.

There is a final point worth mentioning before turning to the nature of the training sets. We need to consider the inability of neural networks to explain their conclusions. Being non-symbolic, neural networks cannot explain to a lawyer why they concluded as they did. All of their intelligence is in the weightings on the links between neurodes. This information is not of a type amenable to symbolic manipulation and hence explanation. The lawyer asking the network why it came to the conclusion it did is going to be very disappointed when its response, like a truculent child, is 'Because.' Therefore, we need to consider how to supplement neural networks' explanatory functions, perhaps by integrating them with rule based systems or by deriving symbolic information from them.

[Bench-Capon, 1993; Bochereau *et al.*, 1991]suggest analysing the weights of the neural network links in order to derive symbolic information. That is, determine which links are most important in the neural network, and from this information draw tentative rules from the network. Quite apart from the difficulty of performing such an analysis, particularly on a large network, one must ask why they are using an inherently sub-symbolic approach to generate symbolic information. If this is all we seek, then we would be better served relying on statistics *simpliciter* or on logical deduction.

Nonetheless, whatever the concerns raised by the nature of neural networks, they are nothing compared with the problems associated with the training sets used to train these neural networks. It is to this issue that we now must turn.

2.2 The nature of the training set

The training set comprises the basis for the intelligence of any neural network, unlike symbolic systems where the encoded rules provide the intelligence. Thus, if we are to generate anything of value from the neural network, we must be careful in choosing the information we encode in the training set. Unfortunately, this has been perhaps the greatest failing on the part of the legal neural network implementors.

There are two basic features of the training set which we need to consider – the types of cases in the training set, and the inputs chosen. Let us look at the cases which the systems to date have chosen to implement.

Types of cases

The first point to bear in mind is the sheer number of cases which neural networks require in order to train properly. This is due to their statistical basis—it is impossible to adjudge any feature as statistically significant unless it is seen in a vast range of cases. Thus, a neural network needs thousands, or at least hundreds, of cases to learn properly. Have we seen this in the legal implementations? With the exception of [Bench-Capon, 1993] and possibly [Bochereau *et al.*, 1991] the systems implemented seem to rely on few cases. For example [Philipps, 1991] relies on ten cases, while [Hobson & Slee, 1994] rely on twenty-six. How can one draw any statistically significant inferences from this small training set? [Philipps, 1991, p. 996] addresses this question. His analysis is based on the notion that there are 'prototypical' cases which define the subject matter. He argues that if one uses only, or mostly, these prototypical cases then one can train a neural network to generate the correct answer. This is nothing more than constraining the training set in order to have the neural network mimic a symbolic reasoner.

To illustrate the problems with this type of approach, let us say that we seek to create a neural network to assess whether a driver will lose her licence to drive because she was drunk while driving. It has two inputs (`drive` and `drunk`) and one output (`licence_loss`). Let us say we can train the network using only two prototypical cases:

Case 1: The driver was drunk and driving and lost her licence.
Case 2: The driver was not drunk and did not lose her licence.

If we train the network with sufficient repetitions, it will generate the expected answer. However, if we are relying only on prototypical cases, that is the existing doctrinal rules, why use a neural network when we could use a production rule system (`IF drunk AND driving THEN licence_loss`) or propositional calculus (`drunk & drive -> licence_loss`)? Both of the other alternatives are more computationally efficient, and at least we can query both as to why they generated the answers they did. Tiny training sets which assume a correct doctrinal answer seem to me to be a misuse of the neural network.

There is another aspect which keeps appearing in legal implementations – the use of hypothetical cases. Due largely, it seems, to the need for large training sets and the difficulty of obtaining such large sets in law, implementors have chosen to supplement their training sets by hypothetical cases. Some implementations add hypothetical cases to genuine ones, for example [Hobson & Slee, 1994], while others rely only entirely on hypotheticals, for example [Walker *et al.*, 1991; Bench-Capon, 1993]. The distinction is immaterial.

'Padding' the training set with hypotheticals seems at first benign, until we consider that these cases are derived from a rule. That is, a rule is specified (for example, `IF drunk AND driving THEN licence_loss`) and then cases are generated in huge profusion which satisfy the rule. The effect of this upon training the network is, once again, to have the network simulate a doctrinal symbolic rule based system. That is, the network should be able to induce the rule from the cases. Once again, we see doctrinal rule positivism creeping into our use of neural networks. But is this such a problem?

The use of leading cases or statutes as the basis for law is not just confined to neural networks. In fact, it has been the basis for all 'black–letter' doctrinal law teaching since law teaching began. To some degree it is also true of legal practice. In this sense, the developers of legal neural networks might well argue that they are only doing as law teachers or lawyers have been doing for years. There are two responses to this rejoinder, one relating to law teaching and the other to law practice. First, legal realism, critical legal studies or postmodern legal theory is steadily gaining ground on canonical teaching, Hence, the days are long gone when one could apply the doctrine without further justification. Secondly, and perhaps more importantly, while doctrinal system might be practically beneficial in symbolic legal expert systems, they are next-to-useless in sub-symbolic legal expert systems. Symbolic legal expert systems can reason with doctrine and deduce consequences from the mixture of rules and facts. They can be used by lawyers to see the rules on which the law is purported based. However, a sub-symbolic system is unable to explain the basis for the decision and certainly cannot do so by reference to 'rules'. Hence, neural networks developers would be better advised seeking not a normative analysis of the law but rather a descriptive analysis. Since the neural networks can generate statistical correlations very quickly and generate predictions from these correlations, they are excellent at modelling the law as it is, rather than as the judges say it is. Neural networks are well placed to give both practitioners and legal theoreticians accurate predictions about the likely outcome of a case, rather than simply replicating a doctrinal normative analysis.

The blind adherence to doctrine in neural networks is not limited to the use of hypothetical cases in the training set. For when one looks at the genuine cases used in a system such as [Hobson & Slee, 1994] one sees that they use the leading cases of the domain. Since leading cases are exceptional (else they would not go on appeal) using them as the basis for statistical analysis is virtually guaranteed to generate poor conclusions. Should we not then try to avoid implementing in a neural network that disease which [Frank, 1930; Frank, 1949] memorably diagnosed – 'upper-courtitus'? It is my submission that using neural networks means that we must choose a domain where the descriptive power of the paradigm can be used. These domains will be where there are large corpora of similar cases, and they are likely to be found in the lowest level courts of first instance. These domains, like car accidents, marital dissolutions and work related injuries, are much more likely to give us the basis for meaningful network than upper court areas like theft, murder, and so on.

Finally, while examining the choice of cases in the training set, consideration must be given to conflict. In particular, the issue arises: how does a neural network deal with conflicting cases in the training set. Traditionally, in building symbolic systems, the way to deal with conflicting cases was to choose one case as better representing the doctrine, and discard the cases which conflicted with it. One can do this with neural networks, but like its symbolic cousin, this approach is almost completely unsatisfactory. What use is a system which cannot tolerate conflict, since this is such a vital and distinguishing feature of law?

[Philipps, 1991, pp. 992–993] suggests that neural networks can handle conflict in the training set. He says of his neural network that:

> "Remarkably, the network also tolerates contradicting learning patterns. The equilibrium attained by the units will be a compromise between the patterns."

Philipps identifies the means by which neural networks handle contradiction: contradictory cases simply lower the weightings on some of the links. So for example, if we have a number of cases which indicate a positive outcome, and one case which indicates the contrary, we will see the weight reduced on the link associating the facts with the positive outcome. The problem arises, of course, where we have one case which is the most important or most recent, and which best represents the law as it is progressing, and a series of old, dated cases which indicate what the law once was. We will see the important case overwhelmed by the sheer number of less important ones.

Philipps argues however that this is not a concern, since according to him the most salient aspect is whether equilibrium is achieved. He argues that striving for equilibrium in conflicting cases is both appropriate and something neural networks do well. His rationale for arguing that equilibrium is a basis for justice relies on the symbol of law being the scales. This is hardly a persuasive argument. His view of 'justice as equilibrium' does not accord well with some of the major theories of justice advanced in the last twenty-five years; even theories as diametrically opposed as [Rawls, 1972] and [Nozick, 1974] do not rely on notions as morally neutral as equilibrium. However, it is interesting that even Philipps who argues for equilibrium notes that neural networks find it difficult to create equilibrium when given contradictory data. He notes that the inconsistency creates a range of possible results for the networks, which are based not any quality of the data but which rather which depend on the rate specified for the learning algorithm. Such an artefact of the implementation seems neither to satisfy the community's expectation of 'justice' nor Philipps' more modest requirement of 'equilibrium'. More work must certainly be undertaken to examine how we can resolve concerns in the face of contradictory data.

Types of inputs/factors

We need also be wary of what input neurodes represent. [Hobson & Slee, 1994] chose to include input nodes such as:

> "Has the accused picked wild mushrooms, flowers, fruit or foliage for reward"
> "Has land been appropriated by a trustee, personal representative, attorney, liquidator or otherwise"

And so on.

Much of this is probably drawn from the doctrinal basis of theft, and we need not return to our discussion of the dangers of doctrine. However, these inputs might alternatively be drawn from a single case.[2] That is, "Has the accused picked wild mushrooms, flowers, fruit or foliage not for reward" might well be drawn from the unusual fact situation of one case. If this is so, then the network is merely performing a kind of pattern matching – something at which a neural network is adept, but which is trivial since simple Boolean retrieval could perform the same function. Thus, in designing the inputs we need to give consideration to the process of matching, and guard against a kind of simple one to one comparison.

Intelligent identification of relevant input nodes can however be a real strength. They may provide weight to alternative legal theories about a given domain. For example in death penalty cases, rather than expressing what the judges *say* are the important criteria in assessing whether the death penalty is appropriate (for example, 'violence' or 'previous convictions' and so on) the system might give credence to what we think might be better explanatory features (for example, what are the races of the defendant and victim). We may find that the neural network generates accurate predictions of outcomes, without reference to the doctrinal basis.

We still must be careful. For example, we can use factors to allow us to justify any conclusion. Say that there are a whole slew of murder cases against your side. To overcome the weight of these precedents we might argue that none of them have considered a new factor, say that the murder was committed with salad forks. Of course, if one is a lawyer then this type of argument may be a little difficult, but if one is the judge then it is a cinch. Training a network on the old cases will ignore the new factor, and the new case can be resolved by the new factor alone. Thus, any case can be justified by adding a new factor.

The addition of factors is therefore a problem. But perhaps not a huge one, and certainly not as problematic as the difficulties with the training set described above. There are other problems with factors: I have, for example, ignored the inherently subjective nature of the choice of factors, which will bring into play its own concerns. However, there is little that one can do about this problem. Others may be able to suggest alternatives.

3 Conclusion

This paper has assessed neural networks from the perspective of legal theory. There are, I am sure, equally cogent issues raised by the above implementations from a purely artificial intelligence–connectionist perspective, see for example [Thagard, 1991]. It is not this paper's thesis that neural networks are inappropriate for legal domains. But it does argue that we need to give careful consideration to artificial intelligence and legal

[2] There is no evidence of this, as the paper does not explain clearly the origins of the 'indexing points'. If the comments misrepresent the authors I apologise: I use the example for illustration only.

theory before advocating their use. In particular it is worrying that we are seeing the tacit acceptance of rule positivism again. This seems to ignore the true benefits that a statistically based paradigm such as neural networks can bring to our study of law.

Before ending, a closing remark: one of the difficulties in assessing the validity or otherwise of neural networks is the imprecision of the language used to describe the neural networks. A cynic might wonder why authors are so keen to skate fleetingly over the type of network, the type of learning rule, the nature of the training set (particularly the origin of the cases used in the set) and even the results generated. But perhaps a more charitable response would be to recognise that we are still at an early stage in the development of legal neural networks and we all need time to come to grips with the technology. One can only hope that as the technology becomes more accepted, and more acceptable, we see neural networks implemented using a clearer appreciation of legal theory. Perhaps then we will be looking for law in the right places.

References

J.A. Andersen & E. Rosenfeld, editors. *Neurocomputing: Foundations of research.* MIT Press, Cambridge, 1989.

K.D. Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals.* MIT Press, 1990.

T. Bench-Capon. Neural networks and open texture. In *Proceedings of the Fourth International Conference on AI and Law*, pages 292–297, Amsterdam, June 1993.

R. Birmingham. A study after Cardozo: De Cicco v. Schweizer, noncooperative games and neural computing. *University of Miami Law Review*, 47:121–145, 1992.

L. Bochereau, D. Bourcier and P. Bourgine. Extracting legal knowledge by means of a multilayer neural network application to municipal jurisprudence. In *Proceedings ICAIL-91*, pages 188–296, Oxford, 1991. ACM Press.

S.J. Burton. *An introduction to law and legal reasoning* . Little, Brown and Company, Boston, 1985.

S. Fish. Working on the chain gang: Interpretation in law and literature. *Texas Law Review*, 60:551–567, 1982.

S. Fish. Wrong again. *Texas Law Review*, 62:299–316, 1983.

J. Frank. *Law and the modern mind.* Anchor, Garden City, 1930.

J. Frank. *Courts on trial: Myth and reality in American justice.* Princeton University Press, Princeton, 1949.

H.L.A. Hart. *The Concept of Law.* Clarendon Press, Oxford, 1961.

J.B. Hobson and D. Slee. Rules, cases and networks in a legal domain. *Law, Computers and Artificial Intelligence*, 2(2):199–134, 1993.

J.B. Hobson and D. Slee. Indexing the Theft Act 1968 for case based reasoning and artificial neural networks. In *Proceedings of the 4th National Conference on Law, Computers and AI*, Exeter, 1994.

D. Kennedy. Freedom and constraint in adjudication: A critical phenomenology. *Journal of Legal Education*, 36:518–562, 1986.

D. Kennedy. A semiotics of legal argument. *Syracuse Law Review*, 42:75–115, 1991.

J. Kolodner. *Case based reasoning.* Morgan Kaufmann, Los Altos, 1993.

E.H. Levi. *An introduction to legal reasoning.* Chicago University Press, Chicago, 1949.

G.F. Luger and W.A. Stubblefield. *Artificial intelligence: Structures and strategies for complex problem solving.* Benjamin Cummings, Redwood City, 2nd Edition, 1993.

J. McClelland, D. Rumelhart and the PDP Research Group. *Parallel distributed processing: Explorations in the microstructure of cognition.* MIT Press, Cambridge, 1988.

C. Mackinnon. *Towards a feminist theory of the state.* Harvard University Press, Cambridge, 1989.

R. S. Michalski, J. G. Carbonell, and T. M. Mitchell. *Machine Learning: an Artificial Intelligence Approach.* Tioga Press, Palo Alto, 1983.

V. Mital and L. Johnson. *Advanced information systems for lawyers.* Chapman Hall, London, 1992.

R. Nozick. *Anarchy, the state and Utopia.* Harvard University Press, Cambridge, 1974.

G.J. van Opdorp and R.F. Walker. A neural network approach to open texture. In H.W.K. Kasperen and A. Oskamp, editors, *Amongst Friends in Computers and Law*, Kluwer Law and Taxation, Deventer, Holland, 1990.

G.J. van Opdorp, R.F. Walker, J.A. Schrickx, C. Groendijk and P.H. van den Berg. Networks at work: a connectionist approach to non-deductive legal reasoning. In *Proceedings of the Third International Conference on Artificial Intelligence and Law*, ACM Press, Oxford, 1991, 278–287.

L. Philipps. Are legal decisions based on the application of rules or prototype recognition? Legal science on the way to neural networks. In *Atti Preliminari del III Convegno Internazionale Logica, Informatica, Diritto*, Firenze, 1989.

L. Philipps. Distribution of damages in car accidents through the use of neural networks. *Cardozo Law Review*, 13:987–1000, 1991.

J. Rawls. *A theory of justice.* Harvard University Press, Cambridge, 1972.

D. E. Rose and R. K. Belew. Legal information retrieval: a hybrid approach. In *Proceedings of the International Conference on AI and Law*, Northeastern University, Boston, MA, 1989.

D. E. Rose and R. K. Belew. A connectionist and symbolic hybrid for improving legal research. *International Journal of Man-Machine Studies*, 35(1):1–33, 1991.

P. Thagard. Connectionism and legal inference. *Cardozo Law Review*, 13:1001–1004, 1991.

S. Toulmin. *The uses of argument.* Cambridge University Press, Cambridge, 1958.

R. Walker, A. Oskamp, J. Schrickx, G. van Opdorp, and P. van den Berg. Prolexs: Creating law and order in a heterogeneous domain. *International Journal of Man-Machine Studies*, 35(1):35–67, 1991.

D.R. Warner. Towards a simple law machine. *Jurimetrics*, 29:451–467, 1989.

D.R. Warner. The role of neural networks in law machine development. *Rutgers Computer and Technology Law Journal*, 16:129–144, 1990.

D.R. Warner. A neural network based law machine: Initial Steps. *Rutgers Computer and Technology Law Journal*, 18:51–63, 1992.

D.R. Warner. A neural network-based law machine: The problem of legitimacy. *Law, Computers and Artificial Intelligence*, 2(2):135–147, 1993.

P. H. Winston. *Artificial Intelligence.* Addison-Wesley, Reading, Massachusetts, 3nd edition, 1992.

J.M. Zurada. *Introduction to artificial neural systems.* West Publishing, St Paul, 1992.