

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

T. Bench-Capon , Legal Theory and Legal KBS: A Computer Scientist's Perspective, 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. 33-42, 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 Theory and Legal KBS: A Computer Scientist's Perspective

Trevor Bench-Capon  
*Dept. of Computer Science*  
*University of Liverpool, PO Box 147, Liverpool, L69 3BX, UK*  
*E-Mail: tbc@compsci.liverpool.ac.uk*

## Abstract

In this paper I shall explore the relationship between legal theory and legal knowledge based systems from the standpoint of a computer scientist engaged on the development of legal KBS. First I shall examine some jurisprudentially based criticisms of a particular approach to legal KBS made by Robert Moles. I shall then examine a second jurisprudential approach found in the work of Richard Susskind. Finally I shall attempt to draw some conclusions for the future relationship between the two disciplines.

*Keywords:* Legal KBS, Legal Theory, Rule Based Representation

## 1 Introduction

In this paper I want to discuss the relation between legal knowledge based systems and legal theory, from the perspective of a computer scientist with a long standing interest in knowledge based systems and their application to the legal domain. I shall begin with a fairly detailed examination of some criticisms of a common approach to building KBS in law advanced from a jurisprudential standpoint by Robert Moles [Moles, 1992]. For a fuller defence of this approach see [Bench-Capon, 1994a] and [Bench-Capon, 1994b], from which the third section of this paper is substantially derived. I then consider another jurisprudential investigation, represented by the well known work of Richard Susskind [Susskind, 1987]. Finally I shall attempt to draw some tentative conclusions as to what a fruitful relationship between legal theory and computer science in this area might be.

## 2 An Approach to Legal KBS

In this section I shall describe my overall approach to building legal knowledge based systems. This is by no means the only approach that can be taken, nor do I wish to argue here that it is the best approach. I believe it to be a practical approach, and a legitimate approach, and it is stated here since I wish to defend it in the next section against the jurisprudentially motivated criticisms of Moles.

Essentially I have consistently argued that a knowledge based system in law should be based on an executable formalisation of law. At the core of the system we should find legislation, but this will need to be supplemented by representations of additional

case law and interpretative material. [Bench-Capon, 1989], [Bench-Capon, 1991] and [Bench-Capon & Coenen, 1992] may be taken as reasonably representative.

My favoured form of representation has always been broadly within the tradition of logic programming. This can be seen as giving rise to a 'rule based' representation, and the execution of the system as the deduction of the logical consequences of the represented knowledge. This is, essentially the style of system criticised by Moles [Moles, 1992], and I shall look at his criticisms in the next section.

### **3 A Jurisprudentially Motivated Attack on the Above Approach**

Moles' objections can be summarised as "Law is not rules" and that in consequence any representation of law in a rule based formalism will traduce the law. Moles first made this point not in the context of legal knowledge based systems, but in the context of jurisprudence: his book *Definition and Rule in Legal Theory* [Moles, 1987] is largely a critique of Hart and his view of jurisprudence which suggests that rules do have a crucial role in law [Hart, 1961]. It is indeed possible to see his assault on rule based legal knowledge based systems, first introduced towards the end of [Moles, 1987], as a continuation of his attack on this jurisprudential view. Moles' own position emphasises the importance of "consequentialism" whereby the consequences of a decision are giving over-riding importance in deciding cases. He writes ([Moles, 1987, p. 217]) of

"the complex decision making which is involved in any legal case, ... it involves a process of complex negotiation which can only be carried out in the light of the many considerations obtaining at that time. ... In stating that legal rules can be applied without further judgement; that they apply in an all or nothing fashion; that legal decision making follows the form of the syllogism or that it is a pattern matching routine, the modern positivists, joined now by the computer scientists take us along a dangerous road."

It is not for me to resolve this dispute between jurisprudes. I will, however, make the following comment. Whilst Moles has made out a persuasive case that cases are, on occasion, decided in a way which appears to conflict with previously accepted rules, and with previous decisions, and, moreover, that such cases may *rightly* be so decided, this does not mean that deduction has no role to play: other examples can be cited to show that this is clearly not the case. Any suggestion of a dichotomy is misleading. For a full, and to me convincing, advocacy of the case for deduction in legal reasoning see MacCormick [MacCormick, 1978].

From my experience while working for the Department of Health and Social Security, I can say that when formulating legislation it was always recognised that it would be impossible to foresee every circumstance to which the legislation would be applied, and hence it was not desirable that the legislation be too tightly drafted. That is, the intention was to lay down certain very specific conditions, such that a benefit would not be payable to anyone who was under 16, or over 70, or not resident in the United Kingdom, but that certain other conditions should be specified in such a way that the adjudicator would have a fair degree of room for manoeuvre so as to reach a just decision in the light of actual circumstances. One specific example concerns the now happily defunct benefit of "Housewives Non-Contributory Invalidity Pension". This was a benefit

payable to married women, who had failed to establish a connection with the labour force through the payment of National Insurance contributions, but who were, through reasons of ill health, incapable both of paid employment, and of “a substantial amount of their normal household duties”. The use of the term “substantial” here was to allow just the kind of flexibility of which we have been speaking, and to avoid the impossible task of defining “household duties”. In practice the wording was too loose and there was great variation in the treatment of people, depending on the views of individual adjudicators. When ultimately the matter was determined by a Tribunal of Commissioners issuing a definitive interpretation of the provision it was interpreted rather more liberally than had been intended; and the legislation was accordingly amended, to bring it into line with the original intention of the policy makers. Two points are worth noting: that the legislators intended to allow flexibility to take account of individual circumstances *within* a framework of more rigid conditions, and that when the amount of flexibility was demonstrated to be too great they felt able to alter the situation by further legislation. In this case the interplay between deduction from rigid conditions and judgement as to circumstances seems clear, at least from the perspective of those responsible for drawing up the legislation.

Moving from the nature of law to computer science issues, there are two points which need discussion: first that use of a rule based representation involves a commitment to the belief that the law is, in some sense, a system of rules; and second, that use of a rule based legal expert system requires one to believe that the rules “can be applied without further judgement”, and that they apply “in an all or nothing fashion”. I believe that neither of these statements is true.

### 3.1 Systems Based on Legislation

Suppose first of all that a fairly extreme version of Moles' position was the case: that the judge could decide a case, say a benefit claim, in whatever way seemed right to him, unfettered either by the written legislation, or by the precedents of past cases, and that the onus of establishing the claim lay on the claimant. We could record this state of affairs in the following way:

```
award-benefit-to(Claimant,yes):-adjudicating-claim(Judge,Claimant),
                                seems-deserving-to(Judge,Claimant).

award-benefit-to(Claimant,no).
```

Note that we must identify the judge responsible for making the decision, since the opinion of other judges can carry no weight here. Note too that the burden of proof is placed on the claimant, since if the judge cannot answer affirmatively to the second condition of the first rule the benefit will not be awarded. Here we have a rule based system, even though no rules, other than those deriving from the procedures, are of relevance.

Admittedly the utility of such a system will be small, although it does make clear that the matter is one of free discretion, *for the appropriate judge*. Moreover, note that the second condition can only be determined by the appropriate judge: any other person will be able to produce no more than an informed guess, based on the facts of the case and knowledge of the judge and his past behaviour.

In practice, even where wide discretion is permitted to the judge, there will be certain conditions which the legislators desire to be satisfied. Thus, for example, the benefit may be intended for persons over a certain age, resident in this country, and these conditions may be prescribed in the legislation. The situation will now be:

```
award-benefit-to(Claimant,yes):-adjudicating-claim(Judge,Claimant),
                                age(Claimant, A),
                                A > 65,
                                resident-in(uk,Claimant),
                                seems-deserving-to(Judge,Claimant)

award-benefit-to(Claimant,no).
```

Here we have what is the usual case: a mixture of supposedly clear cut conditions, and a degree of flexibility, although the flexibility is still rather extreme. The effect of the system now is to direct attention towards two specific points: under-age claimants and those living abroad would see little future in claiming, and the judge would need to be satisfied as to the claimants fulfilling these two conditions before exercising discretion. The problem, on the consequentialist view, would occur only when a judge felt that the claimant ought to be awarded the benefit, but was either not old enough, or not resident. Here the judge is still able to decide in favour of the claimant, although the justification of the decision may prove tricky: "resident-in" is something of a term of art, and so gives scope for creativity, but if the claimant was born only 55 years ago, the judge might need to resort to some fairly extreme interpretation, perhaps holding that such a person should be considered as 65 on grounds of health, the toughness of their life, or something similar.

What this shows is that the role of the judge is crucial, and that the rules of the expert system cannot be interpreted in an all or nothing fashion. In effect, the predicates, no matter how clear cut they may seem, require the assent of the appropriate judge, and so the correct formulation should perhaps be:

```
award-benefit-to(Claimant,yes):-adjudicating-claim(Judge,Claimant),
                                holds(Judge, age(Claimant, A)),
                                holds(Judge, A > 65),
                                holds(Judge, resident-in(uk,Claimant)),
                                seems-deserving-to(Judge,Claimant)

award-benefit-to(Claimant,no).
```

One might, however, consider this an unnecessary complication: the need for the appropriate judge to subscribe to the various decisions might be held to be capable of remaining implicit in the representation, to be brought out in the use of the system. Since the system is designed for a particular practical use by trained operators, why should the knowledge base need to be written so as to make all the assumptions explicit? Thus if it were the judge using the system, he would need to bear in mind at all times that he was making a legal decision, with due regard to the consequences, and that the answers to the questions should be given in the light of this. Similarly if it were the advisor

of a claimant using the system, he would need to remember that the judge would take these considerations into account, and so the answers given by the advisor could never be more than informed predictions of the view that the judge would take. So construed, of course, both users recognise that the rules of the system cannot “be applied without further judgement”, and that they do not apply “in an all or nothing fashion”: such users must recognise that the rules do no more than identify issues, which must be decided in a jurisprudentially appropriate fashion.

### **3.2 Adding Expertise**

So far we have considered only a system that might come from formalising legislation alone. To make a truly useful system we will want to incorporate expertise, knowledge of how these conditions are likely to be interpreted. For a general discussion of how we can graft expertise onto a formalisation of legislation see [2]. In the simplest case this may just be the inclusion of some arithmetical ability, so we might add:

```
Age(Person, A):-date-of-birth(Person, Date)
                difference-in-years(Date, Today,A).
```

Such a rule would provide a convenient way of calculating the age of the person, but not the only way: it is not intended to curtail the freedom of a judge to arrive at a different answer, nor to preclude an age being attributed even where the date of birth of the person is unknown. We must, however, distinguish sharply between the failure of the first condition, which indicates that the date of birth is unknown, and the failure of the second condition which indicates that the claimant's date of birth is known to be too recent. For the rule to be evaluated correctly by the user, the reasoning of the system must be transparent. This point will be further discussed below.

We can, moreover, extend the scope of expertise to give guidance even on discretionary predicates. Suppose that widows had always been found deserving of this benefit. We could add a clause:

```
seems-deserving-to(Judge,Claimant):-widow(Claimant).
```

to represent that every judge finds every widow deserving.

A rule such as this must, however, be clearly flagged to the user, so that it is distinguished as a generalisation from case law. It does not supply a binding sufficient condition, but only a generalisation of past experience, the sort of generalisation that experts might be expected to reach by reflecting on their knowledge of past cases, but one which has no real status and may well turn out to be false. A judge will not be obliged make this induction from the past experience, and so is under no obligation to follow this rule. An advisor using the system must be fully aware of the judge's freedom to make some different induction and so to ignore the rule, although it would be helpful to know that this is the typical conclusion.

This places the relation between past cases and a decision in a particular case in an interesting light: the adjudicator must, it seems, make the induction of the rule afresh each time a case is decided. Only when the adjudicator commits to the rule in the context of a case can deductions from that rule be made. This helps to explain why,

when presenting a case, it is not the rule that is put forward, but the decisions, albeit in such a way as to *encourage* the induction of the desired rule by the adjudicator. In passing we should note that this applies whatever the source of the rules, whether they are stated by an expert, or induced by some machine learning algorithm.

If therefore we wish to incorporate such rules in a knowledge based system, it is essential that we present to the user the backing for the rule, indicating that it is a possible induction from a set of cases, rather than being an expression of legislation. In this way the users can commit to, and hence accept the rule, or reject the induction and hence the rule and the conclusions that may be derived from it, as they see fit.

For these reasons it is important that the results of the system must be presented to users in a way which enables them to distinguish the different sources of and authorities for the various rules employed, and to see the backing for, and status of, such rules, so that they can commit themselves to the rules employed and to the base facts which license the application of those rules. The users will therefore require more than a simple proof trace which constitutes the standard "How?" explanation of typical expert systems. Rather the user must be presented with an *argument*, with the steps made clear. Techniques for presenting the output of the system in the form of an argument susceptible to such a critique based on the argument schema of Toulmin [Toulmin, 1958] can be found in my previous work, such as [Bench-Capon *et al.*, 1993].

Thus when examining the output from such a system, the user must be aware that every assent to a predicate, and every commitment to a rule represents a separate legal decision, which must be made in the context of the case under consideration. Where the user is an adjudicator, the user will have the power to make, and the responsibility to make, these decisions. Where the user is an advisor, the user must predict the decisions that the adjudicator will make. Where the user is an advocate, the user must persuade the adjudicator to make these decisions in a way favourable to the client. Thus the various potential users of a system stand in quite different relations to the rules of the system, and must take their relation to these rules into account in their use of the system.

### 3.3 Summary of Moles' Objections

Throughout the above discussion there may be some confusion as to the meaning of the word 'rule'. As used in the context of rule based systems it is no more than a licence to infer a particular conclusion. One suspects that it may well mean something different in a jurisprudential context, and that Moles is applying that sense to the computer systems. So in the context of a rule based system we can call a representation of the fact that all unmarried men are bachelors such as:

```
bachelor(X) if male(X) and unmarried(X)
```

a rule. But of course there is no kind of *obligation* on unmarried men to be bachelors. It is a rule in the sense that the system is instructed to conclude `bachelor(X)` from the conjunction of `male(X)` and `unmarried(X)`. The rule is directed towards the *system*, not the *user*, and certainly not at the subject under consideration.

We can now see that the problems that trouble Moles should stem not from the use of rules in the representation as such, but from the use that a rule based system might be put to, if used in a simplistic and unreflective way. The users must be able to understand

their relation to the rules of the system, and different types of rule and their authority must be distinguished. Use of Toulmin's argument schema is one way to make explicit the underlying argument which allows the users to commit appropriately to the various steps.

Moles' arguments thus have power only against naive rule based systems, used naively. His challenges may be fair in the case of the original British Nationality Act system [Serhot *et al.*, 1986], which can be seen as representing a rather extreme view. They are not, however, powerful when directed against the bulk of subsequent work in this tradition which takes a much more subtle view of the role of the user, and nature of the representations, whilst remaining firmly in the rule based tradition.

Properly used, rules are not "all or nothing", nor does building a rule based expert system in law rely on identifying rules which have some pre-existence in a Platonic world of jurisprudence. Rather the rules are justified by the artifact they constitute: a match-stick model of St Paul's is a model of St Paul's, and no criticism on the grounds that St Paul's is made of different materials can be made - unless we misuse the model, to examine the fire resistant properties of St Paul's perhaps. A model has an intended use, and similarly a rule based system must provide support *for a class or for classes of user*, and those users must be aware of several things: of their relation to the system; of the way in which the conclusion of the system is reached; and of the part they must play in enabling the conclusion to be reached. A rule based system thus needs a fair degree of sophistication, both in the presentation of its reasoning and, perhaps more importantly, on the part of its users. Given the required sophistication, a rule based system can provide effective support to a user carrying out a task.

## 4 Susskind's Jurisprudential Enquiries

A rather different use of jurisprudence in relation to legal KBS can be found in the work of Richard Susskind [Susskind, 1987; Susskind, 1989]. Susskind embraces a very different position within legal theory from Moles, and has no problem with rule based representations: indeed both his best known systems are based on this formalism. He is, however, equally convinced of the need for legal theory. Legal theory is important for Susskind in two ways: first it is supposed to guide the legal knowledge engineers in building systems - with the claim that this guidance will produce a better quality system; and second legal theory can "articulate the latent jurisprudential presuppositions, assumptions and implications" ([Susskind, 1987], p25) that must inevitably lie within any legal expert system. I shall consider each of these "jobs for jurisprudence" in turn.

### 4.1 Jurisprudential Contribution to Expert Systems Construction

With regard to the first point [Susskind, 1989] lists eleven specific aspects on which he claimed jurisprudence was of help:

1. confirming the suitability of the domain
2. selecting suitable software
3. structuring the expertise
4. eliciting the academic [as opposed to the experiential] knowledge from the expert
5. individuating the law



6. interpreting legislation and case law
7. identifying the limitations of the system
8. coping with the system's lack of knowledge of its own limitations
9. designing the system's inferencing
10. solving problems of time in relation to legal validity
11. identifying the users, function and scope of the system (p29)

This is an impressive list, but whilst it may have been true that the jurisprudential training of the authors of the Latent Damage System, did provide effective guidance to them in that particular case, it is by no means possible to extend this to say that it is an essential factor in the satisfactory performance of these activities. If we take the last point, identifying the users, function and scope of the system, then we find that Susskind's conclusions are little different from those that emerged from my discussion of Moles' critique above: that is they are sensible conclusions, but ones which do not depend essentially on jurisprudential expertise, but conclusions which could equally well emerge from a pragmatic consideration of the system in use.

The problem is that there are no examples of what difference the jurisprudential approach makes: both Susskind's projects tackled new problems, whereas the benefits of jurisprudence could only have been showed by re-addressing a task previously unguided by these considerations and showing how the jurisprudentially informed system was superior in these respects. For example, it would be interesting to explore how the jurisprudentially informed individuation of the law would differ from the isomorphic approach of [Bench-Capon & Coenen, 1992] which is based only on the surface structure of the written sources, and what *computer science* benefits could be obtained with regard to efficiency, verifiability and maintainability.

This is not to deny that there may be advantages: it is rather to seek clarification of what the advantages are, and justification of the claim that this is the only route by which they can be attained.

## 4.2 Articulating Latent Assumptions

Here the role of jurisprudence seems clearer: clearly it is for the jurist to articulate the implications of the use of a legal KBS. The practice of law logically predates legal theory. Law is a creation of society, and legal practice brings it into existence. Legal theory arises as theorists comment on that practice, to explain the phenomena of law, and perhaps also to prescribe good and bad practice. In the main, however, theorists tend not to be practitioners, and practitioners do not in general give very much thought to legal theory in carrying out their day to day activities, even where their actions are in fact influenced by the prescriptions of the theorists. On this view we argue that those concerned with building knowledge based systems in law should be seen as practitioners - they provide additional raw material which can be accommodated in legal theory, but are not constrained to work in a manner dictated by any theory. This is in agreement with the views of Niblett, an early worker in legal expert systems, who wrote:

“a successful expert system is likely to contribute more to jurisprudence than the other way round ... the value of an expert system will not reside in its conformity to some jurisprudential theory” [Niblett, 1981].

Of course, when the jurists have developed a sufficiently powerful account of enough legal knowledge based systems and the jurisprudential assumptions they embody, they may be in a position to argue about what constitutes good and bad practice in building such systems, and this can then feed back into KBS practice. But it is important that these prescriptions emerge from a consideration of KBS activity, rather than simply being stated from some *a priori* position. Moles does in fact prescribe the interesting approach of Ronald Stamper [Stamper, 1991], but his analysis of, and sympathy for other approaches is insufficient to give much weight to his opinion. I have discussed Stamper's approach elsewhere [Bench-Capon, 1994b].

## 5 Legal Theory and Legal KBS

The relationship between legal theory and legal KBS that has emerged from this discussion is as follows. Legal knowledge based systems are built by computer scientists (or by jurists acting as such), and legal theory can then be used to analyse them, laying bare their implicit jurisprudential presuppositions, assumptions and implications. This analysis will, of course, need to be carried out by a legal theoretician who has a good understanding of, and preferably some sympathy for, the systems under analysis. Sometimes, of course, implementor and analyst will be the same person, or members of the same team.

From this analysis some views as to what is good practice, and what systems are acceptable, may emerge: but this is for the jurist - the role of the computer scientist is to produce feasible and clear examples that can form the raw material of the analysis.

For a really fruitful analysis - one which moves from opinion to a more scientifically based assessment - controlled experiments will be needed. As jurisprudential guidelines emerge, perhaps existing systems should be re-implemented following these guidelines, and the gains from following them made explicit. Ideally these gains should be in terms of specific features meaningful in computer science terms, such as improvements in validation, usability or maintainability.

Of course, jurists also disagree amongst themselves. Where this is the case there would be much to be said for them expressing their disagreements in a form which would lead to different prescriptions for computer models. The implementation of the different models that result from the different prescriptions could then form part of the jurisprudential argument, providing a concrete focus for the comparison of the rival claims.

## References

- Bench-Capon, T.J.M. Deep Models, Normative Reasoning and Legal Expert Systems, *Proceedings of 2nd International Conference on AI and the Law*, Vancouver 1989. ACM Press: New York, pages 37-45.
- Bench-Capon, T.J.M. Practical Legal Expert Systems: the Relation Between a Formalisation of Law and Expert Knowledge. In Bennun and Narayanan (eds) *Computers, Law and AI*, Ablex: Norwood, NJ, 1991, pages 191-201.
- Bench-Capon, T.J.M., and Coenen, F.P. Isomorphism and Legal Knowledge Based Systems, *AI and Law*, Vol 1, No 1, 1992, pages 65-86.

- Bench-Capon, T.J.M., Coenen, F.P., and Orton, P. Argument Based Explanation of the British Nationality Act as a Logic Program. *Computers, Law and AI*, vol 2 No 1, 1993, pages 53-66.
- Bench-Capon, T.J.M. In Defence of Rule Based Representations for Legal Knowledge Based Systems, *Proceedings of the 4th National Conference on Computers, Law and AI*, Exeter 1994, Exeter University Centre for Legal Interdisciplinary Development, pp 1-11.
- Bench-Capon, T.J.M. Legal Knowledge Representation: What can be done with rules?”, *THINK Quarterly*, vol.3, june 1994, Institute for Language Technology and AI (ITK), Tilburg University, pages 54-61.
- Hart, H.L.A. *The Concept of Law*, Clarendon Press: Oxford, 1961.
- MacCormick, N. *Legal Reasoning and Legal Theory*, Clarendon Press: Oxford, 1978.
- Moles, R.N. *Definition and Rule in Legal Theory*, Basil Blackwell: Oxford, 1987.
- Moles, R.N. Expert systems - The Need for Theory, in C.A.F.M. Grutters, J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and V.N.J. de Vey Mestdagh (eds), *Legal Knowledge based Systems: Information Technology and Law*, JURIX '92, Koninklijke Vermaede, Lelystad, 1992 pp113-22.
- Niblett, B. Expert Systems for Lawyers. *Computers and Law*, 29, 1981.
- Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., Cory, H.T. The British Nationality Act as a logic program, *Communications of the ACM* 29, 5, 1986, pages 370-386.
- Stamper R.K. The Role of Semantics in Legal Expert Systems and Legal Expert Reasoning, *Ratio Juris*, Vol 4 No 2 pages 219-244.
- Susskind, R.E., *Expert Systems in Law*, Clarendon Press: Oxford, 1987.
- Susskind, R.E. The Latent Damage System: A Jurisprudential Analysis. In *Proceedings of the Second International Conference on AI and Law*, Vancouver, 1989. ACM Press: New York, pages 23-32.
- Toulmin, S. *The Uses of Argument*, Cambridge University Press: Cambridge, 1958.

## **Acknowledgement**

The arguments against the position of Moles in the third section of this paper are largely taken from [Bench-Capon, 1994b].