Legal knowledge based systems
JURIX 92
Information Technology and Law

The Foundation for Legal Knowledge Systems
Editors:
C.A.F.M. Grütters
J.A.P.J. Breuker

H.J. van den Herik
A.H.J. Schmidt
C.N.J. de Vey Mestdagh

T.J.M Bench-Capon, Support for Policy Makers: Prospects for Knowledge Based Systems, in: C.A.F.M. Grütters, J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt, C.N.J. de Vey Mestdagh (eds.), Legal knowledge based systems JURIX 92: Information Technology and Law, Lelystad: Koninklijke Vermande, 1992, pp. 41-49, ISBN 90 5458 031 3.

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

 http://jurix.bsk.utwente.nl/

SUPPORT FOR POLICY MAKERS: PROSPECTS FOR KNOWLEDGE BASED SYSTEMS

T.J.M. BENCH-CAPON
Department of Computer Science, The University of Liverpool, Liverpool, England

Summary

This paper discusses the potential for providing knowledge based support for the task of formulating policy, and determining what legislation is required to implement the policy. From a discussion of previous work in this area, certain major obstacles are identified. Chief among these is the need to match what the KBS can do with the way in which policy makers conceptualise and perform their task. Effective support can only be provided by a system which can be fully integrated into the working practice of its users. Some examples of an alternative approach, based on hypertext, are discussed, and some proposals for overcoming the obstacles with a combination of the hypertext and knowledge based approaches are given.

1. Introduction

There are very many examples of the application of knowledge based systems techniques to the area of law. The overwhelming majority of these have concentrated on answering questions as to how the law applies in a particular case, whether for the purposes of adjudication or of advising a party to the case. There are, however, other tasks related to law which do not consider individual cases, and for which it has been thought possible to supply knowledge base support. One of these tasks is drafting - a task identified long ago in Allen [1957] - and another is the task which must take place even before the legislation can be drafted, namely the formulation of policy, whereby politicians and their advisors decide what laws there should be. It is this last task that will form the subject of this paper.

An early and sustained attempt to address the problem of formulating legislation was carried out as part of the Alvey DHSS Demonstrator project [Bench-Capon, 1987][Taylor & Bench-Capon, 1991][Storrs, 1991]. More recent efforts to address this task by knowledge based means include Debrock et al. [1991], Svensson et al. [1991], and Breuker & Den Haan [1991].

This previous work will be reviewed in the second section of the paper. In the third section some conclusions will be drawn, so as to identify the distinctive issues that arise in this area, and which have to date prevented full success. In the fourth section several systems which offer support through hypertext will be examined. In the final section a new direction for knowledge based support of policy formulation will be described, which will combine some of the better features of the hypertext approaches with some of the better features of the knowledge based approaches to give some promise of providing effective, intelligent, support.

2. Previous systems to support policy formulation

In this section we will describe some previous attempts to address the policy formulation problem.

2.1. The Alvey DHSS demonstrator prototypes

The Alvey DHSS Demonstrator project examined three tasks relating to Welfare Benefits in the United Kingdom: adjudication of claims, advice to members of the public and the formulation of legislation. The fact that these three different areas were addressed gave invaluable insight into the differences between systems in the different areas, and the extent to which knowledge could be shared across applications. The formulation task resulted in a series of four prototypes [Storrs, 1991]. The first prototype [Bench-Capon, 1987] was based on a logical model of legislation which could be queried by the policy maker at various stages throughout the life cycle of a policy problem. A six stage description of the life cycle of a policy problem was used, the stages being:

1. Verify that the problem is real: often complaints are received which are based on a misunderstanding of the provisions of the legislation, or on some difference in the goals of the complainant and those of the government. If the complaint is of this nature, everything is working as the policy maker desires, and there is no problem to solve.
2. Characterise the problem precisely: if there is a genuine problem the first thing to do is to see what the problem is, and to explain why it occurs. A complaint may be originally couched in terms of, say, disabled sixteen to eighteen years olds, where as the problem may actually apply to all members of this age group who are neither in full time education or working. A precise understanding of the problem and its underlying mechanism is an essential prerequisite of solving it.
3. Modify outdated theories: often problems are caused by the underlying theory of society used to formulate policy becoming outmoded. Many problems arise in UK Welfare Benefits because they were devised in a social context, particularly with regard to women, which no longer applies. Many aspects of the UK legislation make sense only if women are seen as dependents of their husbands: if this was true in the 1930s it is no longer generally acceptable.
4. Generate candidate solutions: given a problem there are often a number of different ways to solve it, each of which will have advantages and disadvantages. Successful policy making depends on considering all the available options.
5. Evaluate candidate solutions: here the advantages and disadvantages of the various options must be listed and compared, and a decision, essentially political in nature, taken as to what is the most desirable solution.
6. Specify solution for legal draftsmen: the policy must now be turned into law by specialist draftsmen. They will require precise instructions, so that they fully understand the policy they are trying to effect.

Like all life cycle models, this description is something of an over simplification: there is not really a clean division between the stages, nor a steady progression through them. Indeed it can be said that policy problems are never solved, but only temporarily resolved. None the less it is a useful tool for seeing the range of activities that need to be supported.

The nature of the logical model used in this prototype deserves some consideration here, since it is necessarily very different from that which can be used in an adjudication system. In an adjudication system one has a certain number of facts, and one wishes to draw a particular conclusion of the basis of them. The policy task, however, does not deal with individuals, but with classes of individuals: thus the facts are not available, and reasoning must proceed at a higher level of abstraction. Moreover the direction of the reasoning is not given at the outset: it may be equally necessary to prove either that all pensioners are over a certain age, or that all people over a certain age are pensioners. Thus the logical model must be more flexible and capable of working from abstracted concepts rather than only with the "leaves" of the notional proof tree.

The first prototype confirmed that the model could provide support for the various stages of the policy process, but it was found impossible for policy makers to use it. The

interface of this prototype was very crude, questions and answers being couched in terms of the logical model and presented in a scrolling results window. The second prototype was an attempt to improve the interface, using a desktop metaphor with folders and special "stationery" for asking queries and modifying rules etc. This did not, however, greatly improve matters. As in the case of the first prototype the use of the model and its output could not be appreciated by policy makers.

In order to explore the question of what would be an acceptable interface, the third prototype was developed in a radically different form. Policy makers see their task as constructing arguments: arguments why things are right or wrong, arguments as to why a solution will solve a problem, and arguments as to why one solution should be preferred to another. The third prototype was thus designed as an "argument" processor, in which policy makers could structure their arguments by entering free text into a system giving a graphical representation of the argument based on the argument schema of Toulmin [1958]. The arguments could be divided into sub-arguments, commented on and annotated by colleagues, and rationalised into a final form. This tool was understood, and well liked, by policy makers, but contained no knowledge based component.

The final prototype was intended to link the argument processor to a logical model so that arguments could be verified and filled out by reference to the model. Unfortunately, partly because of lack of time, and partly because of several unresolved conceptual difficulties, notably concerning modality, discussed below, and the mapping from proof to argument, discussed and resolved in Bench-Capon et al. [1991], the integration was not achieved, and the argument processor and logical model ran in the same system but with little interaction.

2.2. A system to compare social security laws of several countries

The system described by Debrock et al, is designed to explore differences in the social security legislation applicable in various member countries of the European Community. The first stage was to examine the legal effects, with later plans to determine and evaluate the micro- and macro-economic consequences of the differences. The technique was to model the various legislative provisions in a form executable by ADS (Aion Development Shell), so that differences in the conditions triggered could be noted. For example it can discover that the relative sex of the two members of a co-habiting couple matters in the Belgian system, but is unimportant in the Dutch system.

One important point of the representation is that it is done at a high level of abstraction: because the intention is not to advise on particular cases, only the main patterns of the regulations need be considered. This is consistent with what was discovered on the Demonstrator project: policy makers can, and do, think in broad terms, unconcerned with how the law may be operationalised and interpreted by those who must apply it in particular cases.

This interesting system is thus directed at the support of a particular policy task, one which is of increasing importance as movement towards harmonisation of Community legislation occurs. It could also be adapted to deal with several alternative proposed systems within a single country. Its focus is, however, rather specific: it makes no attempt to tackle the policy process as a whole, nor to tie the legislation to the policy aims of the policy makers. This in turn avoids many of the problems of policy support, but perhaps it may be sensible to provide a dedicated tool with a defined scope.

2.3. ExpertiSZe

The ExpertiSZe system [Svensson et al., 1991], is designed to be used by legislators who are in the process of preparing new legislation. The basic idea is to model the

legislation in a knowledge based system, and then to use this model to simulate the effect of the legislation. In particular the simulation is intended to answer questions as to the completeness of the legislation - are all cases covered? - and the consistency of the legislation - is there a single outcome for every combination of circumstances? It will also identify prescriptions which play no part in the outcome of any case. The simulation, in conjunction with a database of sample cases, designed to reflect the distribution of features across the population, will also predict the effects of the legislation which can then be examined for consistency with policy goals and constraints.

Such a system has much in common with the logical model approach of the first two Demonstrator prototypes, and one suspects that while the benefits of using the model are real, there may be similar difficulties involved in introducing it into the working pattern of policy makers. One point about the representation must be made; the model is very similar to a model that could be used by an adjudication system, and operates by simulation. Thus the reasoning requires a set of fully specified descriptions of circumstances, whether determined from the model or from a sample database. This limits the functionality, since the facts determined by the possession of a legal attribute cannot be derived from the model, and it may not be possible for policy makers to work at the level of abstraction that they wish. None the less it is a further interesting example of the way in which a logical model can be used to support policy formulation.

2.4. TRACS

The Tracs system [Breuker & Den Haan, 1991], similarly operates by simulating the operation of a model of the regulations, in this case regulations applicable to road traffic law. A key feature of this system is a clear separation of the regulation knowledge base from a world knowledge base, which contains a representation of relevant common sense and a number of scenarios which can be used to generate the descriptions of situations to which the regulations will be applied to determine their effect. The use of the model would be to examine the simulated effects of the legislation to determine what is permitted and what is forbidden by the regulations so that it can be seen whether the policy goals are achieved or not.

Both the Tracs system and ExpertiSZe can supply valuable input to certain stages of the policy process. They do, however, suffer from a problem inevitable upon using a simulation based approach, namely an inability to reason about the model itself, which places the burden of the analysis of the results of simulation on the policy maker. Questions of the form: "what would need to be the case for this goal to be achieved?", a very typical policy question, can only be answered by reasoning about the model, not by executing it.

3. Problems for the above approaches

The above systems demonstrate the potential of support for policy formulation, but this potential will only be realised if the systems can be made usable by policy makers, and can be presented in a form which can be accommodated in the work patterns of policy makers. This may be of less importance where the system is directed to a specified, well defined, stage of the process, but this compartmentalisation will necessarily limit the usefulness of the system, and may make the overhead required to construct the logical model unacceptable. The ideal is pervasive use of the system throughout the policy process, with the formation and amendment of the model an integral part of that process. But there is strong evidence, particularly from the demonstrator project, that current knowledge based techniques cannot be so integrated and made acceptable to policy makers. It is worth here trying to explain why the interface of the first two Demonstrator prototypes could not be made acceptable.

3.1. The policy makers view and the KBS view - a conceptual clash

The reasons for the inability of policy makers to relate to the logical model were that access to the model was through a set of very low level tools for examining and retrieving information from the logical model representing the knowledge pertinent to the problem which formed the Knowledge Base (KB). Interaction with it, both in terms of the selection and asking of queries, and of the output received in response to queries, was couched entirely in terms of the KB. Thus use of the model and the interpretation of the output was possible only to someone who understood in detail, and thought in terms of, the KB and who was additionally able to interpret the questions a policy maker might wish to pose in terms of the KB. Consequences were:

1. The policy maker had no clear expectations as to what would be produced by the various queries, and was unable to choose the query appropriate to his problems;
2. The policy maker was unable to understand the output: this was partly because it was very voluminous and partly because it often contained, in addition to the interesting consequences, irrelevant and sometimes trivial consequences. Whilst someone familiar with the knowledge base was able to see these as not relevant almost immediately, this was not so for the policy maker;
3. The policy maker was unable to relate what was produced by the logical model to his problem;
4. The above problems were compounded by the need, in order to answer a policy question, to link several queries together, in a way determined by the output from a previous query.

In brief the first two prototypes were unusable because there was no clear understanding of their relation to the problems under examination, and because neither the tactics nor the strategy required to use it could be understood except in terms of the underlying logical model of the KB.

These problems run too deep to be tackled by cosmetic changes to the interface, such as, for example, simply presenting the individual sentences returned from the model in a more English-like form. What is needed is a rethinking of what will be produced as output from the logical model to ensure that it is problem-orientated, rather than KB-orientated. This is essential if the system is to integrate in a task where the users see themselves as weighing issues and constructing arguments. These problems will apply equally to any system which conceptualises in terms of a KB. A possible way out of this dilemma will be discussed in section 5.

3.2. Abstraction and system orientation

If we can make a system acceptable, it will be important to consider what will be special about the knowledge represented in a policy formulation system. In contrast to systems designed to apply the law, it will be necessary that the system be able to operate with concepts that are highly abstracted; the detail that surrounds individual cases cannot be considered, and plays little role in the thinking of policy makers. Thus a policy maker may decide to introduce a supplement payable to those on benefit who live in homes which are "difficult to heat". The policy maker may be perfectly content to work with this vague concept, and leave the operational definition and interpretation of the concept to those responsible for applying the law. When making policy, and when drafting legislation, in may even be undesirable to go into any more detail. Thus the model must be able to handle these abstracted concepts, and must not rely on detailed information to be executable.

Moreover it is important that the system be able to reason so as to discover properties of the model, rather than simply relationships between descriptions of circumstances and outcomes. It must be possible to find out why conclusions cannot be drawn, and what

would license, or block, conclusions. Further it must be possible to reason from a conclusion to the circumstances that would license the conclusion as well as from circumstances to conclusions. This means that we must build the model using a flexible and manipulable representation unbiased to any particular use. This places a high level of demand both on the representation, and on the program that will manipulate it [Bench-Capon & Forder, 1991].

3.3. Modality

One other point we must consider here is the pervasive use of modal qualifiers when discussing policy. Toulmin's original schema for arguments [Toulmin, 1958] contained an explicit place for recording such qualifications. In the third Demonstrator prototype this element was omitted, but was continually re-introduced by policy makers entering the text of their arguments. The basic question a policy maker answers is "What should be done?". Note that this is not the should of deontic logic, policy makers are not discussing ethics and morality, but a pragmatic "should": - "what is the most desirable course of action?". Whilst this kind of modality is inescapable in reasoning about policy, the logic of such qualifiers is utterly unclear. We can say, perhaps that if P will result in Q and Q is undesirable then P is prima facie undesirable. But notice that all solutions may be prima facie undesirable, and if we have a choice between three exhaustive undesirable options, then the least undesirable becomes desirable.

There are problems here both as to what needs to be implemented, and how it might be implemented. But without some answer to these problems the reasoning of the model will remain peripheral to the policy maker's concerns and mode of thought, where thinking in these terms is fundamental.

4. Hypertext approaches

Given the difficulties in providing a support environment using KBS techniques, we should consider some alternative approaches here. Hypertext is becoming increasingly fashionable, and it is unsurprising that it has been applied to the policy problem. The idea here is fundamentally similar to that enshrined in the third Demonstrator prototype: that useful support can be provided by a system which does not attempt to reason, but which enables policy makers to marshal their arguments in a structured fashion.

Most notable of the hypertext based systems for policy support are those based on the IBIS (Issue Based Information Systems) method developed by Hans Rittel [Kunz & Rittel, 1970]. The idea is that policy problems can be seen as groups of issues, on which people take positions, which they support by arguments. The goal of IBIS is to enable the users to understand the conflicting positions, identify the key points of contention and structure their thoughts as they strive for consensus. The imposed structure helps to keep the discussion constructive, and to focus on what is agreed and what remains to be resolved. At the conclusion of the discussion all the material to provide a reasoned case for the selected option is available.

Such a system with its range of node types, issues, positions and arguments, and clearly defined links between them is ideal to implement as a hypertext system. Examples of implemented systems are the gIBIS system [Conklin & Begeman, 1988] and the hyperIBIS system [Isenmann, 1992]. Both of these systems are fairly straightforward hypertext implementations of the IBIS method, providing a means of creating and manipulating nodes and links to record a discussion in the appropriate form. The purpose of the tool is only to structure the discussion: all reasoning and intelligence must come from the users. There are, for example, no checks as to whether an argument really does support a position. In some ways this can be an advantage: the users of gIBIS were pleased with the lack of enforcement which placed no constraints on what they could

express, whilst making it easier they felt to detect inconsistencies in arguments that might have slipped through unnoticed in standard prose.

Conklin and Begeman in particular have a very full discussion of the strengths and weaknesses of their system as shown in practical use. It is worth quoting their conclusion:

"Our experiences suggest that the computer is indeed a powerful medium for collaboration and debate among members of a team, but that the integration of computers into the fine detail of real work is attended by some severe breakdowns. Some of these breakdowns are due to inadequate interfaces, others to inappropriate underlying representations, and still others to insufficiently rich models of work practices and methods." [Conklin & Begeman, 1988, p. 330]

Thus it is interesting to see that integrating even such a tool - apparently based around the working practices of its users, and relatively unconstraining - has problems of integration with the task. This, particularly in view of the areas where these problems arise, suggests powerfully that to attempt such an integration with an unadorned logical model will be doomed to inevitable failure.

5. Some proposals

So where does this leave the potential for knowledge based support of the policy process? It seems clear that if we are to offer any kind of effective support it is essential to integrate the systems smoothly into the working practices of policy makers; otherwise the systems will not be used. This can be done by constraining the function of the proposed systems so that they address only an individual, clearly specified task, such as checking the proposed legislation for consistency. On such a view the role of KBS is limited. But is it possible to do more?

I believe that a solution may be possible on the following lines. At the highest level, where issues are imprecisely specified, and where the problems lack structure, the hypertext based approach has many attractions. It allows thoughts to be stated and explored without constraint, and mirrors the way policy makers work at present. But as the arguments become more detailed, precision becomes necessary, and it is here that knowledge based techniques may have a role. In the IBIS like systems for example, arguments are advanced in support of positions. Whether the argument in fact does support the position, and whether it coheres with other arguments for the position is a question of logic, not opinion. And this logic can be tested by reference to a knowledge based model. To analogise: it would be impossible to hold all discussions in terms of syllogisms, because the form constrains expression too much. But once we have agreed on the desired conclusion and the available premises, expression of the argument as a syllogism is an effective demonstration of its soundness. So too the constrained but rigorous reasoning supplied by a logical model can be used effectively if it is located in the proper context.

The proposal is therefore that the system be constructed around a hypertext facility, but that at appropriate points, when the issues are sufficiently clear and the questions can be posed at a sufficient level of precision, access should be given to an underlying model. At this point the policy maker can be clear as to what is required from the KBS and what the information returned from the KBS means. Importantly also, at this level the questions of modality are avoided; these questions are resolved at a higher level, in the freer form discussion. The KBS can demonstrate that certain consequences flow from certain proposed actions: that these consequences, and hence the actions are undesirable is a matter for statement and resolution at the higher level.

There remain, however, problems of presentation of the output from the KBS. Logical proofs, the natural form of output from a KBS, are notoriously uncongenial to policy makers. Even here they will require the organisation and structure that transforms a proof into an argument. (See Bench-Capon et al. [1991] for a full discussion of the distinctions between a proof and an argument.) Such access to the KBS must therefore be couched in terms of arguments, and Toulmin's scheme, whilst too constraining for the higher level discussion, offers an acceptable structure for such arguments. The improvement in explanation that can be provided by a legal KBS by using Toulmin's structure for explanation, and a mechanism for so presenting the explanations is presented in Bench-Capon et al. [1992a].

Straightforward explanation is not, however, enough to meet all the needs of policy makers. Not all the questions which can be posed by policy makers to a KBS take the form of a proposition to be proved. Sometimes too, the policy maker will want to contribute extra information, both factual and in the form of generalisations to the debate. What is really needed, therefore, is the ability to carry on a dialectical dialogue with the KBS. Only in this way can the necessary parity of status between policy maker and system, and the essentially exploratory nature of the work be achieved. The creation of such dialogues, the consequential improvement in interaction, and an architecture for their implementation is described in Bench-Capon et al. [1992b].

We have therefore a cluster of techniques which can be grafted onto the bottom of a hypertext tool for policy discussion which will supply the rigour and reasoning capabilities of a knowledge based model at a level where they can be used and appreciated by policy makers, while staying within the consistent framework provided by the adoption of argumentation as the central notion. In this way it may be possible to exploit the potential benefits of KBS for policy formulation.

6. Conclusion

Much of this paper has been concerned with the identification of barriers to providing effective support for policy makers. It is not enough to produce technical solutions: the policy makers must be able to relate to these solutions. Some suggested lines of attack have been described. Conklin and Begeman end their description of gIBIS with the words:

"we are just at the beginning of a long but exciting path, which will culminate when we have succeeded in making tools as effective and transparent in structuring communication as the telephone has grown to be in simply transmitting it." [Conklin & Begeman, 1988, p. 331]

With these words I concur: I hope also to have shown how knowledge based techniques can lie on that path.

7. References

This article may only be cited as: Bench-Capon, T.J.M., Support for Policy Makers: Prospects for Knowledge Based Systems. In: Grütters, C.A.F.M., J.A.P.J. Breuker, H.J. van den Herik, A.H.J. Schmidt and C.N.J. de Vey Mestdagh (eds), *Legal Knowledge Based Systems: Information Technology & Law, JURIX'92*, Koninklijke Vermande, Lelystad, NL, 1992.

[Allen, 1957] Allen, L.E., Symbolic Logic: A Razor-Edged Tool for Drafting and Interpreting Legislation. In: *Yale Law Journal*, no. 66, 1957.

[Bench-Capon, 1987] Bench-Capon, T.J.M., Support for Policy Makers: Formulating Legislation with the Aid of Logical Models. In: *Proceedings of the First International Conference on AI and Law*, Boston, 1987.

[Bench-Capon & Forder, 1991] Bench-Capon, T.J.M. and J.M. Forder, Knowledge Representation for Legal Applications. In: Bench-Capon, T.J.M. (ed), *Knowledge based Systems and Legal Applications*, Academic Press, 1991.

[Bench-Capon et al., 1991] Bench-Capon, T.J.M., D. Lowes and A.M. McEnery, Argument-based Explanation of Logic Programs. In: *Knowledge Based Systems*, vol. 4, no. 3, September 1991.

[Bench-Capon et al., 1992a] Bench-Capon, T.J.M., F.P. Coenen and P. Orton, Argument Based Explanation of the British Nationality Act as a Logic Program. In: *Proceedings of the Third National Conference on Law Computers and Artificial Intelligence*, Aberystwyth, 1992.

[Bench-Capon et al., 1992b] Bench-Capon, T.J.M., P.E. Dunne and P.H. Leng, A Dialogue Game For Dialectical Interaction with Expert Systems. In: *Proceedings of the 12th International Conference on Expert Systems and Their Applications*, Avignon, 1992.

[Conklin & Begeman, 1988] Conklin, J. and M.L. Begeman, gIBIS: A Hypertext Tool for Exploratory Policy Discussion. In: *ACM Transactions on Office Information Systems*, vol. 6, no. 4, October 1988.

[Breuker & Den Haan, 1991] Breuker, J.A.P.J. and N. den Haan, Separating World and Regulation Knowledge: Where is the Logic? In: *Proceedings of the Third International Conference on AI and Law*, Oxford, 1991.

[Debrock et al., 1991] Debrock, K., V. Lemmens, F. Robben and B. van Buggenhout, Development of a Knowledge Based System for the Comparison of National Social Security Systems of Member States of the European Community in a Dynamic Perspective. In: Noortwijk, C. van, A.H.J. Schmidt and R.G.F. Winkels (eds), *Legal Knowledge Based Systems: Aims for Research and Development, JUIRX'90*, Koninklijke Vermande, Lelystad, 1991.

[Isenmann, 1992] Isenmann, S., HyperIBIS - A Tool for Argumentative Problem Solving. In: *Proceedings of the Third Conference on Database and Expert Systems Applications (DEXA'92)*, Valencia, 1992.

[Kunz & Rittel, 1970] Kunz, W. and H.W.J. Rittel, *Issues as Elements of Information Systems, Working Paper No. 131*, Institute of Urban and Regional Development, University of California, Berkeley, 1970.

[Storrs, 1991] Storrs, G., The Policy System. In: Bench-Capon, T.J.M. (ed), *Knowledge based Systems and Legal Applications*, Academic Press, 1991.

[Svensson et al., 1991] Svensson, J., P.J.M. Kordelaar, J.G.J. Wassink and G.J. van 't Eind, *Analysing Legislation with ExpertiSZe*, The Hague Conference of the European Group of Public Information, 1991 (not published).

[Taylor & Bench-Capon, 1991] Taylor, A.D. and T.J.M. Bench-Capon, Support for the Formulation of Legislation. In: Bench-Capon, T.J.M. (ed), *Knowledge based Systems and Legal Applications*, Academic Press, 1991.

[Toulmin, 1958] Toulmin, S., *The Uses of Argument*, Cambridge University Press, 1958.