

DEEP MODELS, NORMATIVE REASONING AND LEGAL EXPERT SYSTEMS

T.J.M.Bench-Capon
Department of Computer Science
University of Liverpool
Liverpool
England

Abstract

This paper discusses the role of deep models and deontic logic in legal expert systems. Whilst much research work insists on the importance of both these features, legal expert systems are being built using shallow models and no more propositional logic, and are claimed to be successful in use. There is then a prima facie conflict between findings of research and commercial practice, which this paper explores, and attempts to explain.

Introduction

There is a growing fashion in expert systems to prefer systems which are said to have deep knowledge over those which have only shallow knowledge. Whilst there is no firm agreement on what is meant by "deep knowledge", (but see Bobrow 1984 for some typical examples), shallow systems are usually characterised as those which are based on heuristic rules representing expertise elicited from an expert in the domain, whereas deep knowledge systems have an underlying model which reflects the structure of the domain. The relation between the deep model and the set of shallow rules is basically that the shallow rules could be derived from the deep model by reasoning appropriately from first principles. Proponents of the use of deep systems claim that they will provide a number of advantages. Among the more commonly cited are that they will enable easier knowledge acquisition, more robust systems, more maintainable systems, systems which have re-usable knowledge bases, and systems with improved explanation capability. The case for systems with deep knowledge in the field of legal expert systems has been put forward by, for example, Thorne McCarty. In (McCarty 1984) he advocated using a deep conceptual model to support legal expert systems, saying

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0-89791-322-1/89/0600/0037 \$1.50

"the most critical task for the development of an intelligent legal information system ... is the construction of a conceptual model for the relevant legal domain."

The desire for deep models is common to all expert systems domains, but with regard to legal expert systems, there is a second source of dissatisfaction with the shallow rules that form the basis of common expert systems. Laws do, amongst other things, confer rights and impose duties, and set out the boundaries of what is (legally) permitted and (legally) obligatory. In other words, laws express norms. This has been taken to imply that a legal expert system has to be a normative system, that it needs to be capable of embodying normative reasoning. This has further led some people to insist that a deontic logic is required to underpin such systems. The argument is well put in (Susskind 1987, p225):

"characteristically we find in our law-formulations the use of typically normative vocabulary such as must, ought, may, should, shall, and variants thereof. How indeed might knowledge of law-formulations be manipulated while reasoning other than by the application of some deontic logic (logic of normative concepts)? The necessity for the inclusion of some deontic logic within a legal inference engine would seem beyond dispute"

McCarty also sees the necessity of a deontic logic. His "Permissions and Obligations" (McCarty 1983) represents one of the more fully thought out deontic logics developed for use in legal expert systems, and has focussed attention on this aspect of legal expert systems. (McCarty, it should be said, does not, however, see the deontic modalities as being unique in their applicability to law: he also stresses the impor-

tance of other modalities, such as the epistemic and temporal modalities.)

McCarty is a particularly instructive example because his views that a conceptual model and a deontic logic are required to enable legal expert systems to overcome certain limitations inherent in shallow systems were developed in the course of coming up against these limitations in the course of producing the earliest versions of TAXMAN (McCarty 1977), and the voice of experience always deserves a hearing. However, these views are in conflict with other notable views in the field, such as those of expressed by the Logic Programming Group at Imperial College, in, for example, (Sergot 1985). They argue, against both of the above views, that deontic logic and normative reasoning is unnecessary in a legal expert system, and that a deep conceptual model can be rendered unnecessary if a formalisation of the relevant legislation is used as the basis of a such an expert system.

If we look beyond academic systems to the current work being done in legal expert systems, especially that which has a view to commercial exploitation, we find that it is largely based on shallow systems using at most first order predicate calculus as the underlying logic, and pays little attention to ensuring that the knowledge base is isomorphic with respect to the legislation from which it derives. Two examples of such systems may be mentioned here (and will be discussed later): the Retirement Pension Forecast and Advice System (RPFA) developed for use within the UK Department of Social Security, described in (Springel-Sinclair and Trevena 1988) and the Latent Damage Advisor (LDA) developed by Richard Susskind and Philip Capper, described in (Capper and Susskind 1988). Both of these systems were implemented using expert system shells offering no more than propositional logic (the Aion Development System for RPFA and Crystal for LDA). The former is now a fully operational system, which has paid back its development cost in less than a year of use. The success of such systems presents us with an apparent difficulty in that they pay scant regard to the prescriptions emerging from the research community. These prescriptions therefore require examination to show whether deep models are essential, whether deontic logic is necessary, and whether faithfully representing the legislation is a worthwhile activity. The intention of this paper is to provide just such an examination. In particular I shall try to uncover the element of truth in all three of the positions, by trying to show when and why deep models and normative reasoning are required, so as to identify a class of systems in which they can indeed be dispensed with, and by sketching the advantages that derive from using a representation

of the legislation.

Shallow Models

The paradigmatic examples of shallow systems tend to be derived from a variety of fault diagnosis systems, where the fault may be in the human body (as with the classic of the genre, MYCIN), a washing machine or a car. Such systems are meant to capture the expertise of a person experienced in the task, and the rules often resemble tips that one might receive from an old hand teaching one the task. Thus an experienced mechanic might tell a novice that if a car will not start it is a good idea to see if the headlights are working, as, if they are not, the probable cause of the malfunction is a flat battery. Such a nugget of wisdom, and others like it, could be made the basis of production rules, or a logic program, and a shallow expert system constructed to execute them. What makes the model shallow is not that it is expressed in rules, but that it contains no information as to the various causal chains which underlie the empirical associations, noticed in the course of performing the task and encapsulated in the rules. The structure of the domain is essentially causal, whereas the rules do not express causal relationships. In the battery case, for example, the form of the rule might lead one to suppose suppose that the malfunction in the headlights caused the flatness of the battery in some way whereas the reverse is, of course, the case. The expert may or may not be aware of the causal chains; it makes no difference to his ability to apply the rule. Sometimes too, the empirical association may have no causal basis at all; the favourite of the amateur TV repairman is that if all else fails, give the set a good kick and that gets it going again. Sometimes the faith in the empirical associations may be no more than superstition.

In fact, of course, especially if our expert is reliable, the empirical associations will be firmly grounded in implicit causal chains, and the rules and resulting system may well be quite effective at diagnosis. In the case of a diagnosis system this is all we ask.

But now consider what a system in the legal area built on the basis of empirical associations derived from an experienced person would be like. Take as an example a system to determine whether people were citizens of a particular country. Anyone who spends much time at airports and international conferences can try to identify the nationality of people in the throng milling around him, and will probably be able to develop a few crude empirical associations, in terms of general appearance, accent, dress and bearing, which will enable him to achieve a reasonable degree of accuracy. Perhaps immigration officials

whose days are spent looking at people and their passports could have even better empirical associations and a very high degree of accuracy. We could encode such knowledge in an expert system and use it to give an opinion on citizenship. But, no matter how accurate the answers it gave, it would be inconceivable to use such a system to apply the law, and it would be stretching a point to call such a system a "legal expert system". For there is in fact no reference to the law at all. While the empirical associations employ terms, like accent and appearance, which are linked in a complex and indeterminate way with the things, like place of birth and parentage, which form the legal basis of citizenship, they themselves have no legal force, and could not form the basis of an argument which could be presented in law. While a diagnosis system can perhaps accept being judged purely by the accuracy of its conclusions, a legal expert system needs to show that these conclusions have been arrived at in a valid way.

Legal expert systems thus have an important difference from diagnosis systems in the role assigned to explanation. Diagnosis systems require explanations so as to allow confidence to be invested in their conclusions. The explanations of legal expert systems in contrast are in many ways as important as the decision itself, for they represent the (legally crucial) argument which supports the decision. Thus whilst the diagnosis may be acceptable independently of the explanation, and the explanation need be no more than persuasive, the relationship between decision and explanation in a legal system is closer in that the decision is valueless without the argument which leads to it, and that argument must itself conform to certain conditions in order to be an acceptable argument. A legal expert system must not only arrive at the correct decision, but must arrive at it in the correct way.

This explains why no one tries to build a legal expert system in the above style. The knowledge elicited from an expert in the legal domain is not simply a matter of empirical associations. Rather a good deal of the relevant knowledge will be the understanding of the law that the expert has derived from his reading of the legislation and his consideration of case law. It is the marshalling of this diverse material into a form where he can apply it to a case that corresponds to the development of empirical associations in the diagnostic domains. Both RPPFA and LDA, for example, were based on rules supplied by experts in their respective fields: in RPPFA these were derived mostly from statute, and in LDA mostly from case law. In both cases however the source forms had undergone a considerable degree of processing to yield the form of the rule. Empirical associations are not entirely missing from

the legal field, however, and examples may be found in (Bench-Capon 1989a). None the less, any rule in such a system, should be ultimately capable of justification by reference to some piece or pieces of legislation and/or some case or cases so that an acceptable legal argument can be given. But even with systems where all the rules can be justified with the chapter and verse of a precedent or paragraph, such systems remain shallow because they have no straightforward linkage to the statute, or to the other rules making up the system.

Systems based on a Formalisation of legislation

A legal expert system based on a formalisation of the relevant written law, in contrast, does have an explicit linkage back to the statute, since every rule in such a system is supposed to be a formal paraphrase of some clause of the legislation. For examples of such systems see (Sergot 1986) and (Bench-Capon 1987). In practice, it should be said, such formalisations fall short of complete isomorphism with the legislation they formalise. There are still problematic constructs used in legislation which require the development of general techniques to guide their formalisation. For a discussion of one such construction, the counterfactual conditional, see (Bench-Capon 1989b). Isomorphism remains, however, the aspiration, and the more faithful the formalisation, the greater the advantages of the approach. In what follows I shall assume that the formalisations can be taken as being entirely faithful.

Although systems of this kind have some of the superficial characteristics of shallow systems, most notably in that they consist of a set of rules, they are, in fact, quite different. The rules do not record empirical associations generated by experience, or summaries of knowledge gleaned from diverse sources, but represent definitions of legal terms which owe their existence to the piece of legislation formalised. Such definitions are either in terms of other terms defined elsewhere in the legislation, or in terms not so defined: in the latter case the user must be asked whether or not they succeed, or the formalisation must be augmented in some way. For a fuller discussion of this point see (Bench-Capon 1988a).

This does represent a significant gain over systems based solely on knowledge elicited from an expert. Let us review the claims made for systems based on deep models. First it was said that they will enable easier knowledge acquisition; this is fulfilled in that the formalisation of a fragment of law is demonstrably simpler than teasing knowledge out of an expert. Second it was said that they will produce more robust systems, in the sense that unenvisaged cases are less likely to occur; unlike the opinion of an expert, which

will inevitably be determined by particular features of his experience, a formalisation is designed, as is the legislation, to cover all cases, without reference to their frequency. Third it was said that we would get more maintainable systems; given a faithful formalisation we ought to be able simply to delete superseded laws and replace them with the appropriate amendments. Fourth it was argued that deep knowledge systems would have re-usable knowledge bases; since there is no task dependant knowledge in the legislation, a knowledge base founded on a formalisation will be equally applicable to any task using that legislation. Lastly it was hoped to improve our explanation capability. Again an explanation in terms of the statute constitutes an acceptable legal argument. So we can see that a legal expert system based on a formalisation of the relevant legislation fulfils many of the hopes of a deep knowledge system. For a fuller discussion of these points see (Bench-Capon 1988b). Why then can we not regard such a system as a deep knowledge system, and say that the formalisation is the deep model, just as the circuit diagram is the deep model for a fault diagnosis system of the circuit?

McCarty said that a deep model was meant to supply

"a language in which we can describe legal cases."

A formalisation, by being a set of definitions of terms specific to such a language, is, in a real sense, such a language. But the formalisation approach does run up against limitations, and we need to consider those now, for we shall see that whilst we can describe cases using this language, it will not permit us to say everything we need to say about them.

Limits of formalisations

The limits of the formalisation are encountered because one will inevitably arrive at a term not itself defined in the legislation, and hence also undefined in the formalisation. Such terms must be applied by the user to a particular case in the light of the facts of the case, his understanding of language, such precedents as may exist, general legal principles and his knowledge of the world to which the legislation refers. If we are lucky, such terms will be straightforwardly understood by the user, and he may therefore find little difficulty in applying them. This was the case with the British Nationality Act project (Sergot 1986), where we arrived at questions turning on the age of the applicant, his place of birth, his father's citizenship, and the like. But in other cases we will get obscure terms and terms redolent with vagueness, as for example, in UK Social Security Law a claim made after a prescribed time limit will not be accepted

unless "good cause" for its lateness can be demonstrated. Here the user will know that he must decide whether or not the case he is considering does or does not exhibit good cause, and that once he has made his decision the consequences of that decision can be computed by his system. But a system based solely on a formalisation of the legislation can offer no help with this decision. Such help might be offered if we were to augment the formalisation by additional rules which try to break down such concepts into terms that the user can apply. Such decompositions will depend on the user for whom the system is intended; age will require no elaboration for a person trying to apply a citizenship system to himself, but an official making the legally binding decision will need to know that he needs to see a birth certificate before accepting the age of an applicant as verified. If no birth certificate is forthcoming, we will have a "hard case".

These elaborations may be derived from case law, or from official publications produced to provide guidelines to those administering the law, or they may simply be nuggets of experience from some expert. Whatever their source, however, they will never offer more than sufficient conditions for the application of the terms they elaborate. There will always remain the possibility of some case arising whose circumstances have not been envisaged by those doing the elaboration, and in such a case the user will be forced to supply an answer from his own reasoning resources.

Let us consider the case of late claim in a little more detail. Under UK law a claim for benefit has to be submitted within a certain time unless the claimant is able to prove good cause for the delay. We can determine from case law that being confined to one's own house by sickness does constitute good cause, unless one is living with someone who it would be reasonable to expect to forward one's claim on one's behalf. Further examination of case law would show certain categories of relationship for whom this expectation would be reasonable. Now all of this would be a good basis for a shallow model to decide the question, although it would leave gaps. Case law may show a 12 year old child may reasonably be expected to post the claim, whereas a three old child may not, but we may be faced with a six year old. The shallow model breaks down here, and we may wish we had a deep model. But what would a deep model be of? Clearly there is a relationship between the age of a child and the responsibilities with which he can be trusted. But we need to have a pretty complex model of human development, and perhaps model other things such as the distance to the post box, before we could hope to use such a model to decide the question. It may well be thought that the construction of such a model

would present insurmountable difficulties. None of these matters is touched on by the legislation, nor are the factors to consider well bounded. Some of the factors that need to be considered can be inferred from a consideration of factors given weight in past cases, but there is no assurance that this fact does not contain a novel, though relevant factor. The key thing is that a deep model on such a case would not be of the legislation, nor of concepts defined to in the legislation, but rather of the world about which the legislation legislates, which relates to the legislation through terms not defined in, and often not even referred to by, that legislation. Thus the language represented by the formalisation does not allow us to describe the case in terms of all the factors that may be relevant to the decision, but only in terms found in legislation. So we can say that a case exhibits "good cause", but we cannot give our reasons for saying that it does so, without using terms not found in legislation.

Why do we want deep models?

Now let us return to the situation of an adjudicator attempting to decide a case with the aid of a formalisation of law, augmented to include some elaborations derived from case law. He will have his attention directed towards a number of questions and can be assured that the legislation will be applied so as to give the correct result once he has answered these questions. These questions will fall into three categories. First the question may be one of ascertainable fact; thus he may be asked whether the person in question is present in Great Britain. This he can easily and unequivocally discover. Second the question may be one the answer to which follows inexorably from some matter of ascertainable fact. Thus the age of a person who possesses a birth certificate follows without question from observations that can be made on the basis of that certificate. The idea of following without question is partly a matter of knowledge of the world, but is likely to be re-inforced by previous decisions, or some directing guidance from the adjudicator's superiors. The third kind of question is more difficult, however. This is the "good cause" kind of case, where decisions and guidance can do no more than give a number of sufficient conditions for answering the question in an affirmative or negative way, and where, unlike the case of the birth certificate, it is rare for one case to fit exactly with the circumstances of a previous case. So while he may be able to apply law which has already been established, it is often the case that he cannot, because the law is (as yet) silent on the particular point in question. It is in such cases that legal reasoning comes to the fore, since we are no longer dealing with the almost

mechanical application of predetermined rules the applicability of which is not in doubt, but with decisions as to the applicability of a rule to a given case, and where we may have several rules which potentially apply, and which would lead to different conclusions, and where it could be that none of the rules are applicable, and so the decision can only be made according to more general principles.

The situation is that to decide a case where the applicability of the rules is in doubt requires knowledge that goes beyond what can be found in legislation or even in case law. Systems which simply execute rules derived from legislation and case law operate at a syntactic level, and the kind of reasoning required here will often depend crucially on the semantic features of the terms involved. And this is where deep models of the domain come in. At this point it is not enough to know the definition based relationships between the terms used in the legislation, nor sets of sufficient conditions determining the extensions of terms not so defined, but in addition the semantic implications of those terms. To return to the good cause example, it is necessary to model in some way the influence of increasing age on the responsibility which it would be reasonable to assign to a person, and the influence other features would have, such as whether the person in question was living in an urban or rural area. And here a deep model is necessary, if we are to support the reasoning.

At this point it may be objected that all legal reasoning that we would wish to support falls into the last case, and an expert system which could do no more than apply established law would be of little benefit to anyone.

Such a picture may well emerge from a consideration of proceedings that come before a court since in a case of straightforward application there would be no case to argue. But this would be a mistaken picture with regard to the general potential for expert systems in law. Firstly many people do seek legal advice when they have, in fact, no case, because they are ignorant of the law. A straightforward expert system could make their position plain to them. Also, and more importantly in terms of the number of legal decisions made, many of the legal decisions taken in the course of the administration of law in such fields as welfare benefits, tax and employment, do involve the routine application of law, often by people who have no legal qualifications. In such fields it is only in rare or novel cases that lawyers, let alone courts, are involved at all. And, also importantly from the point of view of the exploitation of expert systems, many companies have internal regulations which govern the actions of their employees in a quasi-legal fashion.

The expenses claims of employees of large firms are processed in a way similar to claims for welfare benefits, and similar systems would be expected to support both. Although the regulations which govern such schemes are internal to the company rather than the laws of the land, this makes little difference to the clerk responsible for their application.

The RPPFA system fell into just this class. The system was designed to support clerks advising people (by letter) as to whether they could expect to receive a Retirement Pension when they attained the appropriate age, and if so, at what rate they could expect to be paid. This advice requires calculations based on the employment history and contributions record of the individual concerned, and other factors such as their marital history, children and absences abroad. Naturally this would differ from individual to individual. Further complications exist since certain options are available to some people: thus a married woman may substitute her husband's contribution record for her own in certain circumstances. The manual system suffered from a number of defects, most notably the accuracy of the advice, the time taken to produce it, the cost of producing it, and the quality of the letter sent out. These defects arose not because some cases were hard, in the sense that ground breaking decisions needed to be made, but because of the complexity of the rules, the multiplicity of their interactions, and the fact that some rarely applicable rules tended to be overlooked. There were, however, no problems in determining whether particular rules were applicable in a given case, nor in determining the applicable facts. It was therefore possible to devise a set of rules, not isomorphic with, but founded on legislation, which could handle all the routine cases with confidence. There remained, however, a number of cases which were non-routine, because the applicant had worked for a period abroad. No attempt was made to automate these cases; and the presence of a straightforward criterion to separate the easy cases for those requiring detailed examination meant that there was little degradation in the utility of the system. The system is determinate and requires no judgement; it could have been built in a 4GL, or in a conventional imperative language. The expert systems route was chosen, however, because it provided a speedier implementation (important since savings are in the order of \$1.7 million per annum), and because it was felt to be more maintainable. The RPPFA system therefore provides an excellent example of a system for routine application of the law to decide cases where there are a large number of routine cases, and where these can be easily identified. Note that it is the large volume of cases that need to be decided (currently 300,000+ a

year) that makes this system viable.

My contention therefore is that the RPPFA is not unique, but that there is a substantial class of routine cases which are often entrusted to fairly low level adjudicators, and which can be decided by a fairly uninspired application of the consequences of definitions to be found in legislation and guidance. Only when a difficult case arises need one go beyond this, and only then does a need for a deep model of the concepts underlying this legislation and guidance arise. In the RPPFA situation these could be identified and give individual (human) consideration. More generally, such cases are well exemplified by the kind of precedent setting cases decided by High Level courts and seen as landmark decisions, considered by McCarty in his TAXMAN project (and by others, such as Rissland (e.g. Rissland and Ashley 1987), whose primary interests lie in the nature of legal reasoning). Many people, however, are interested in providing systems to support the routine decision making of low level adjudicators. Such systems can get by with shallow models, although they might find it advisable to incorporate a model of the legislation to provide the sorts of advantages in terms of maintenance and so forth that I outlined above.

To summarise, deep models are advocated in expert systems generally because they provide a number of advantages in the form of maintenance etc.. These advantages can be obtained within a legal expert system by the use of a formalisation, which can, without stretching the point unduly, be seen as providing a deep model of the law. On the other hand the law can be seen as not providing an answer to certain kinds of non-routine case until a decision is made on that particular case. Here no model based solely on legislation could be adequate, because such a model is incomplete until filled out by the decisions. Here the reasoning ceases to be solely about the law, and starts to need concepts drawn from the domain which is the subject matter of the law. Clearly no model of the law is going to provide this, and so a deep model of the underlying domain is required if such reasoning is to be supported. Notice that we are not obliged to support such reasoning; we can always leave such matters to the unaided user, and support only the application of the law once decisions as to the contentious issues have been made.

The kind of system that would result from this approach can be exemplified by the Latent Damage Adviser. The law on Latent Damage is complex and rarely encountered, except by a few specialists. It is, however, closely related to the law on negligence. The system acts as a guide through the law pertinent to Latent Damage, to be used by a competent lawyer

familiar with the law on negligence. The target user group means that the system need not deal with hard cases: where recourse to a deep model would be needed the onus is placed upon the user. The user can see what questions he must answer, and what the implications of his answers will be, but must do all the specifically legal reasoning himself. None the less the advantages of the system are still considerable. The claim is made (Susskind 1988, p 192):

"It is estimated that a competent lawyer would take five to ten hours to understand the Act and its implications. ... By using the Latent Damage System, however, the lawyer can find solutions to latent damage problems in about five to ten minutes. The system is an "intelligent assistant" - it guides a user through all and only those legal rules which bear on the problem at hand."

If anything this understates the usefulness of the system: the expert in question, Philip Capper, an acknowledged expert in the field, and an author of the first book on the statute, found himself surprised on occasions by the output of the system, and that the system was, on reflection, right. This suggests that five to ten hours work would not result in an understanding as good as that provided by using the LDA.

Normative Reasoning

We can now turn our attention to the need for normative reasoning within legal expert systems. Because legislation often uses words which seem to suggest that it is concerned with deontic notions, it has been suggested that a formalisation of that legislation should of necessity employ a deontic logic. For a discussion, see (Herrestad 1988). Whilst it is impossible to deny the presence of these deontic concepts in legislation, it is worth remarking on the extent to which the adjudicator is obliged to use deontic reasoning.

The peculiarities of the deontic notions lie in their interplay with other notions of necessity. Consider the statements "two + two must equal four", "a bachelor must be unmarried", "an unsupported stone must fall", "you must not kill people" "you must register the birth of a child" and "you must wear a topcoat when it is raining". All of these express necessities, but they are necessities of different kinds: the first two are logical or metaphysical necessities, the third a physical necessity, the fourth a moral necessity, the fifth is a legal necessity and the last a mere prudential necessity. Puzzles arise because whereas we can conclude from "it (metaphysically) must be the case that p", that p is

indeed the case, we cannot infer from "it (morally) must be the case that p" that p is the case, since we do not live in an ideal world. So that an obligatory act has many characteristics of a necessary act, it has the important difference in that it need not be done - provided the agent concerned is willing to break the law. But similar situations arise with physical necessities: unsupported stones only fall as long as the natural laws are obeyed. The point is that these notions of necessity apply within certain boundaries: physical necessities are necessary within a certain set of natural laws, and moral necessities are necessary within a certain set of moral laws. The modalities only become important when we conceive of stepping outside the set of laws, and deontic modalities are interesting because of the ease with which we can conceive moral laws being infringed. Provided, however, it is understood that the discussion is within a given framework of laws, be they physical, moral or legal, the relevant necessity can be taken as necessity without qualification.

Now the low level adjudicator in a routine case does not have the option either to make new law, nor to go against existing law, except in the sense that he can cease to do his job as specified and make an eccentric decision which would, it is to be hoped, be overturned on appeal. Thus such an adjudicator recognises these limitations on his freedom of action and should be anxious to make decisions in which what is said to be obligatory by the law is taken as being necessary by him in his adjudicatory role. Such an adjudicator is therefore taking no normative decision beyond one that says that he should abide by the law and perform his adjudicatory task in conformity with the law. If we supply him with a system to support his task, therefore, he should not expect that system to perform any specifically normative reasoning. It may be that it is a normative decision to use the system, but once that decision is made the adjudicator will expect it to reason as though what should (legally) be the case is of necessity the case. Such an adjudicator is working within a given set of legal laws which he has no power to change.

This is the position of the adjudicator whose task is to do no more than apply the law. The position changes when we have a higher level adjudicator confronted with a case in which the law is not clear. Here the adjudicator is forced to resort to more general legal principles, which may be in conflict. He must arrive at a decision, which may in effect be the creation of a new norm if it is to be used as a precedent. Such an adjudicator is going to have to appeal to his sense of justice, both as enshrined in general legal principles and in his vaguer notion of natural justice, to decide

which principles and precedents he is to follow, and what decision he is to make. Such an adjudicator must operate outside the framework provided by the law, because he cannot take the norms expressed in the legal principles as mandatory, since where they are in conflict a contradiction would result. Such an adjudicator must rather do genuine normative reasoning in deciding which to follow. What he is doing is deciding what should (morally or at least extra-legally) be necessary (legally). It is only when this boundary is crossed that deontic reasoning becomes critical.

We have thus drawn a distinction between an adjudicator whose role is to apply the law to routine case to which, it is assumed, the law is readily applicable, and an adjudicator who must recognise that his decision may extend the law because it represents a novel decision. In the former case a knowledge of the law is sufficient, and no normative reasoning need be indulged in. In the latter case, knowledge of the law simply will not enable a decision to be reached, and so a decision must be taken which is justifiable in terms of rather wider notions and which will require normative reasoning.

What we should particularly notice here is that the cases which fall into the latter category are also those where the need for deep models of concepts underlying the law were felt necessary.

Other modalities

The foregoing applies, of course, only to the deontic modalities. Law is concerned with other modalities too; particularly the epistemic and temporal modalities. These modalities present a great challenge to those who would construct legal expert systems by formalising law, since it is not clear that first order predicate calculus is sufficient, or, if it is not, what alternatives should be used. Discussion of these modalities is outside the remit of this paper: one should, however, point out in passing that McCarty expects a deep model to assist in tackling these modalities (and the deontic modalities), and that attempts have been made to accommodate temporal reasoning in a first order framework, e.g. (Kowalski and Sergot 1986).

Conclusion

In conclusion we may point to an important distinction between adjudicators into those whose role is to apply the law as set out in the legislation and the decisions of superior adjudicators, and those superior adjudicators whose decisions become, through the mechanism of precedent, the law to be applied by lower adjudicators. That many adjudicators do fall into the former category should not be doubted; they abound in

government departments administering welfare benefits, tax and other laws. If we extend our notion of law to include such quasi-legislation as the internal regulations of companies which govern such things as the lending policies of banks and the expenses claims of their staff, such adjudicators are found to be very numerous indeed. In their case, it has been argued, neither deep models nor deontic reasoning are necessary in the routine cases which comprise the majority of the cases they are faced with. This means that we can support their task with relatively straightforward legal expert systems. Of course users of such systems will need to accept that these systems are unable to deal with cases which fall outside the routine, and reflect this in their use of the system. In particular an organisation deciding to employ such systems must ensure that effort is made to detect such cases, and to handle them outside of the system. Spotting non-routine cases may be straightforward, as in the case of the RPFA, or it may be difficult, and so require vigilance (and a degree of expertise) on the part of the user so that he can refer them elsewhere, or apply his own talents for modelling the underlying domain and normative reasoning to decide them. None the less there is an enormous potential for such systems.

When we consider the higher adjudicators, however, these simple systems will no longer serve. The cases which reach them do so only because they are outside the routine, and require the deeper understanding that these adjudicators can be expected to bring to bear. Systems to support such adjudicators must therefore both model the underlying domain and have the ability to weigh normative factors.

It is unsurprising that those who have been attracted to the field of AI and Law by a desire to study legal reasoning and who have therefore concentrated their attention on interesting (because non-routine) cases, see a definite need for deep conceptual models and for deontic reasoning. It should, in view of the above discussion, be unsurprising also that others who have concentrated rather on the potential application of legal expert systems for the support (or even replacement) of the routine judgements of low level adjudicators should see such an elaboration as unnecessary. Legal reasoning operates at a number of levels, and the mechanics required to support the highest level is neither necessary nor appropriate to deal with the cases which can be decided by a routine application of law.

References

Bench-Capon 1987. T.J.M. Bench-Capon, G.O Robinson, T.W. Routen and M.J. Sergot, Logic Programming for Large Scale Applications in Law, 1st Interna-

- tional Conference on AI and the Law, Boston, 1987
ACM Press.
- Bench-Capon 1988a. T.J.M. Bench-Capon, Logical Models of Legislation and Expert Systems, in *Expert Systems and Law*, H. Fiedler, F. Haft and T. Traummüller (eds), Attempto Verlag, Tuebingen, 1988.
- Bench-Capon 1988b. T.J.M. Bench-Capon, Applying Legal Expert Systems Techniques: Practical Considerations, in *KBS in Government 1988*, P. Duffin (ed), Blenheim Online Publications, Pinner, 1988.
- Bench-Capon 1989a. T.J.M. Bench-Capon, Practical Legal Expert Systems: the Relationship Between a Formalisation of Law and Expert Knowledge, in M. Bennun and A. Narayanan (eds) *Computers, Law and AI*, Ablex, 1989.
- Bench-Capon 1989b. Representing Counterfactual Conditionals, *Proceedings of Artificial Intelligence and the Simulation of Behaviour*, Brighton 1989, A. Cohn (ed), Pitman, 1989.
- Bobrow 1984. Bobrow, D., (ed) *Qualitative Reasoning About Physical Systems*, published as *Artificial Intelligence Journal*, vol 24, December, 1984.
- Capper and Susskind 1988. P. Capper and R.E. Susskind, *Latent Damage Law - The Expert System*, Butterworths, 1988.
- Herrestad 1988. H. Herrestad, Some problems Concerning the Representation of Legal Norms, 1st Scandinavian Conference on AI, 1988.
- Kowalski and Sergot 1986. R.A Kowalski and M.J.Sergot, A Logic Based Calculus of Events, *New Generation Computing*, Vol 4, pp 67-95, Ohmsha Ltd and Springer-Verlag, 1986.
- McCarty 1977. L.Thorne McCarty, Reflections on TAXMAN: An Experiment in Artificial Intelligence and Legal Reasoning, *Harvard Law Review* 90, 837-93.
- McCarty 1983. L. Thorne McCarty, Permissions and Obligations, *Proceedings of the Eighth IJCAI*, 1983.
- McCarty 1984. L. Thorne McCarty, Intelligent Legal Information Systems: Problems and Prospects, in *Data Processing and the Law*, C.Campbell (ed), Sweet and Maxwell, London 1984.
- Rissland and Ashley 1987. E.L.Rissland and K.D.Ashley, A Case-Based System for Trade Secrets Law, 1st International Conference on AI and the Law, Boston 1987. ACM Press.
- Sergot 1985. M.J. Sergot, Representing Legislation as Logic Programs, in *Machine Intelligence 11*, Oxford University Press, 1985. The position is also expressed in numerous other papers from the group.
- Sergot 1986. M.J. Sergot, H.T. Cory, P. Hammond, R.A. Kowalski, F. Kriwaczek, and F. Sadri, Formalisation of the British Nationality Act. In *Yearbook of Law, Computers and Technology*, Vol 2, C.Arnold (ed), Butterworths, London 1986.
- Springel-Sinclair and Trevena 1988. S. Springel-Sinclair and G. Trevena, The DHSS Retirement Pension Forecast and Advice System, in P.Duffin (ed) *Knowledge Based Systems: Applications in Administrative Government*, Ellis Horwood, Chichester, 1988.
- Susskind 1987. R.E. Susskind, *Expert Systems in Law*, Clarendon Press, Oxford, 1987.
- Susskind 1988. R.E. Susskind, Government Applications of Expert Systems in Law. In P.Duffin (ed) *KBS in Government 88*, Blenheim Online, 1988.