Some Observations on Modelling Case Based Reasoning With Formal Argument Models

T.J.M.Bench-Capon
LIAL- Legal Informatics at Liverpool
Department of Computer Science
The University of Liverpool, Liverpool UK
+151-794-3697

tbc@csc.liv.ac.uk

ABSTRACT

In this paper I shall explore the modelling of case based reasoning using a formal model of argument, taking the approach of Prakken and Sartor as my starting point. I first consider their method of representing cases, and describe how - if we restrict ourselves to independent boolean factors - we can fruitfully model the domain as a partial order on rules. I then consider the issues relating to quantifiable factors, as used in HYPO, and factor hierarchies, as used in CATO. The former presents some difficulties for modelling as a partial order, and, coupled with the latter, forces us to recognise two different kinds of reasoning used in concept application which have different implications for representing the domain. I then present some conclusions arising from the discussion.

Keywords

Argument, case based reasoning, precedent.

1. INTRODUCTION

One of the most fruitful areas of research in Artificial Intelligence and Law has been the modelling of legal argument. There are two major strands to this work. One, which includes the work of Rissland, Ashley and Aleven, and which can be represented in this paper by HYPO (Ashley 1990) and CATO (Aleven 1997), is based on arguing with cases, and is typically applied in domains where precedent is of much importance. Here argument is about identifying some putatively suitable case to follow, and meeting objections to its suitability. Work here tends to be based on a descriptive account of legal reasoning and mainly concerned with emulating (HYPO) or teaching (CATO) the moves made by lawyers when reasoning with cases. The other strand, exemplified by the work of Prakken and Sartor (Prakken and Sartor 1998) and Hage (1997), starts from the need to reason with statutory provisions and exceptions to them, and is concerned to give an account of the logic of legal argument, which can explain the defeasible nature of legal reasoning and its ability to cope with conflicting norms.

ICAIL-99 Oslo, Norway Copyright (c) 1999 ACM 1-58113-165-8/99/6 $5.00

Although the starting points are rather different, case law in one case and statutes in the other, Prakken and Sartor have attempted to apply their logical tools to a reconstruction of the style of reasoning found in HYPO and CATO (Prakken and Sartor 1997 and Prakken and Sartor 1998). In this paper I will take this reconstruction as a starting point, and explore the extent to which it is successful in reconciling the two strands of argumentation modelling.

2. PRAKKEN AND SARTOR'S REPRESENTATION OF CASES

The key idea in Prakken and Sartor's reconstruction is the way they represent precedents. Following HYPO they describe cases by the factors present in the case and the outcome of the case. Factors in a case are classified as pro-plaintiff or pro-defendant according as to whether they favour the plaintiff or the defendant. Now in any given case both pro-plaintiff and pro-defendant factors are likely to be present. Suppose we have a domain with pro-plaintiff factors A,B and C and pro-defendant factors D,E and F. Suppose we have a particular precedent with factors A,B,D,E present, and which was decided in favour of the defendant. Prakken and Sartor represent this precedent as three rules.

One rule expresses the strongest reasons that can be adduced by the plaintiff:

R1: A & B -> p.

Another rule expresses the strongest reasons that can be supplied by the defendant:

R2: D & E -> d.

The final rule encapsulates the decision made in the case, which held that the defendants case was the stronger, by expressing that R2 has priority over R1:

R3: R1 < R2

In this way each precedent will yield three rules. Given the set of precedents represented in this way, they can apply their method, which is based on a dialogue game, of reasoning with conflicting rules and priority relations between them to reason about new cases.

3. PRECEDENTS AS PARTIAL ORDERS

If we adopt Prakken and Sartor's method of representing precedents, and we add the assumption that all the factors are independent of one another[1], and are boolean, either present or absent, we can see that any combination of factors could occur in a case. Thus every combination of factors of a particular tendency (pro-plaintiff or pro-defendant) could potentially form the antecedent of a rule such as R1 or R2. Making the further assumption that adding a factor of the same tendency to the antecedent of a rule makes the rule stronger[2] we can form the possible rules for each tendency into two partial orders, as illustrated in Figure 1.

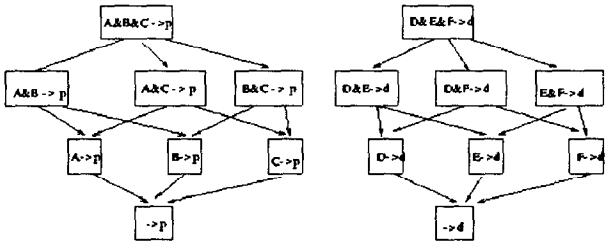

Figure 1: Partial Orders of pro-Plaintiff and pro-Defendant Rules

Now each precedent will supply a priority relation between one of the nodes in the pro-plaintiff partial order and one of the nodes in the pro-defendant partial order, thus joining the two into a single partial order. Suppose for example we have two cases, one containing factors A,B,C,D,E, which was found for the plaintiff, and the other containing factors A,D,E which was found for the defendant. We can now use these two precedents to combine the partial orders as shown in Figure 2.

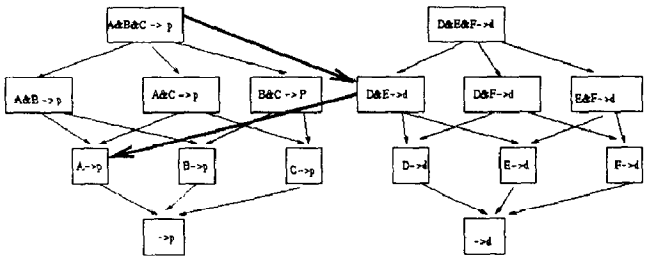

Figure 2: Partial orders with two precedents

Now if we are presented with a new case we can, if the right precedents are available, use this combined partial order to come to a definite conclusion. For example given a case with factors A,B,C,E, we can see that it should be decided in favour of the plaintiff since

```
A&B&C->p > D&E->d > E->d.
```

Sometimes, however, we may not be able to come to such a decision, since the partial order may not prescribe the required priority. For example given a case with factors A,B,D,E, we get rules

```
R4: A & B -> p
```

and

```
R5: D & E -> d
```

and we have no way of determining whether R4 < R5 or R5 < R4. Essentially the case poses the question of whether the absence of C is sufficient to change the decision made in the first case or whether the presence of B is sufficient to change the decision made in the second. Both Prakken and Sartor and the HYPO and CATO systems provide ways of presenting the argument; but neither can resolve the matter (unless we use some notion of burden of proof to ensure that one side must establish their claim or lose)[3].

As more precedents come in we can add links between the two original partial orders and become able to decide more cases. In the example with three pro-plaintiff and three pro-defendant cases there are 64 possible priority relations, of which the two example precedents above allow us to determine seven. In principle we could envisage enough precedents to determine all the priority relations, in which case we could decide any case by means of the partial order – the domain would be sufficiently worked out by case law to remove all hard cases.

In order to be consistent, the partial order must contain no cycles. Suppose for example a third case with factors A,B,C,D was decided for the defendant, giving the situation shown in Figure 3.

---

[1] This assumption is perhaps stronger than necessary. It is important that factors are not exclusive, since that would make some combinations in the partial order impossible, and important that they are not equivalent, since if A is equivalent to B, then A = (A & B) = B. Factors need not, however, be strictly independent, although making this assumption simplifies the analysis.

[2] This assumption is made in HYPO and CATO, but not by Prakken and Sartor In Prakken and Sartor (1998) they provide an interesting example (p 267) where two factors may each favour one side of the question individually, but favour the other side in combination. What follows therefore should not be seen as a reconstruction of Prakken and Sartor's approach, although I do believe that it reflects many of the same intuitions. Also they do provide a scheme of general rules (p 268) to allow the assumption to be made as a (defeasible) default, so that the account in this paper does describe what they regard as the typical case.

[3] This does not really matter in the context of their systems, since their aim is to construct and model arguments that can be made in a case, rather than to reach a decision.

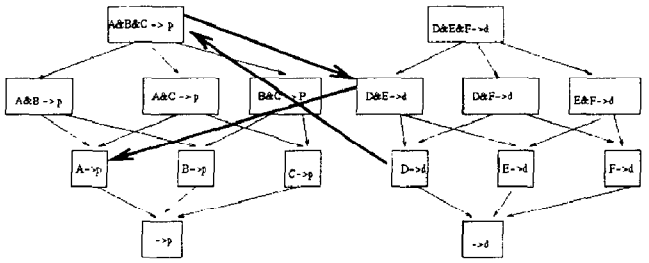

Figure 3:An inconsistent precedent

Now we have a real problem, since we now have irreproachable grounds for saying both that

A & B & C -> p > D & E -> d (precedent 1)

and that

A & B & C -> p < D & E -> d

since

D -> d > A & B & C -> p

from precedent 3, and the pro-defendant partial order tells us that

D & E -> d > D -> d.

The normal method of treating such cycles in partial orders of regarding all three nodes involved in the cycle as equal cannot work: decisions have already been made which say that this is not so. So we must do one of three things:

- assume that there has been some concept drift, in which case we shall want to remove the earlier precedent to break the cycle;
- consider the latest precedent as a rogue decision and ignore it; or
- conclude that our initial analysis was wrong, and there is some fourth pro-defendant factor, present in the latest case and absent in the first case (or some other pro-plaintiff factor present in the first case and absent from the new case) which was sufficient to swing the decision in the defendant's favour.

In the last case we will need to reconstruct one of our partial orders to include this new factor.

Looking at the representation of precedents in this way makes clear many of the features both of HYPO and Prakken and Sartor's approach. Where we can establish a defined ordering between the pro-plaintiff factors and the pro-defendant factors, both approaches will provide an argument to which there is no response. Where there is no such defined ordering we can use the partial orders to present our case. Thus in the terminology of Prakken and Sartor, given the situation shown in Figure 2 and the new case containing factors A,B,D,E, the plaintiff will cite the case containing A,B,C,D, and E, and attempt to broaden R1 to R4, and the defendant will distinguish the case by pointing to the absence of C, and cite the case with A,D and E as a counter-example. In turn the plaintiff will distinguish this case by pointing to the presence of B. Similar citations and responses will be produced by HYPO.

4. DIMENSIONS

In the above discussion factors have been taken to be boolean, as in the examples of Prakken and Sartor. But one of the major insights of HYPO was that many factors are quantifiable, and have a direction, so that a larger value for the factor increases the strength of the case for either the plaintiff or the defendant. Thus in the example of Prakken and Sartor, which concerns change of fiscal domicile, one factor is the duration of absence. Although this seems to be a clearly quantifiable factor, Prakken and Sartor make it boolean by introducing two factors, long-duration (pro-plaintiff) and short-duration (pro-defendant). This allows the factor to be accommodated in our model, although it does violate the assumption and that factors are independent, and so means that some apparently possible priority relations cannot occur. Suppose in our example B was long-duration and F was short-duration; this would mean that no comparison could be made between nodes containing B in the pro-plaintiff partial order and nodes containing F in the pro-defendant partial order, since no case could contain both factors. This is not in itself a problem; in fact it eases the task, since it removes 16 out of the 64 priorities we need to determine.

There is, however, a significant problem in choosing whether to describe a case as containing the factor long-duration rather than short-duration. This means that we need to be able to identify a cut-off point (or points) so that we can classify cases as exhibiting one factor rather than the other. In some cases the cut-off point may be clear: for example the UK law relating to "ordinary residence" requires a period of three years. In others, however, there may be no such clear threshold, and the cut-off point may need to be determined by case law. Indeed, one of the major uses of dimensions in HYPO is to find boundary counter-examples, so as to determine such cut-off points. Argument about where *short-duration* becomes *long-duration* may be the point on which the whole case turns. The problem may be exacerbated where the cut-off point is itself dependent on other factors, as in UK retirement pension where the point at which a person becomes old enough to qualify for the benefit depends on their gender. Such considerations suggest that something is lost by reducing dimensions to booleans.

Prakken and Sartor recognise that HYPO, unlike their proposal, uses and exploits non-boolean factors but comment simply that "It seems to us that there are no theoretical objections to extending our analysis with such features" (Prakken and Sartor 1998, p279). Let us consider whether this is really so by seeing how the introduction of such quantitative factors look in our model as partial orders.

Clearly duration is an issue, present in all cases, but which can, depending on its extent, work in favour of either side. So let us include duration as a factor in both the pro-plaintiff and the pro-defendant partial orders. It cannot, however, be represented by a single node in each of the partial orders, since it varies. Instead we might think that we could represent the duration in a particular case by a node on the arc leading to the node containing the boolean factors which apply, its position on the arc being

determined by how favourable to the plaintiff/defendant it is. Suppose we have two pro-plaintiff boolean factors, A and C and two pro-defendant boolean factors, D and E, and duration. Suppose also we have three precedents, one with A,C,D,E and duration = 12, decided for the plaintiff, one with A,C,D,E and duration = 11, decided for the defendant, and one with A,D and duration = 12 also decided for the defendant. We would represent this as in Figure 4.

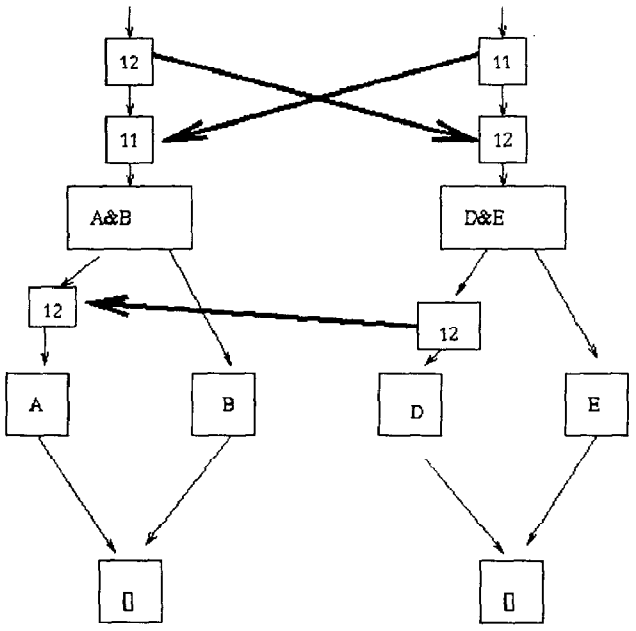

Figure 4: Partial order with a quantitative factor and three precedents

However, this simply will not do. The diagram suggests that the rule

```
A & B & Duration = 6 -> p
```

is stronger than the rule

```
A & Duration = 12 -> p
```

whereas we have no reason to say this. It could even be that `duration > 11` is a necessary condition for the plaintiff to win, irrespective of any other factors. Worse, I can see no better way of representing a quantitative factor in the partial order, unless one constructs a different partial order for every possible duration, which gives rise to intractably many potential priorities. It therefore seems to me that there is a theoretical difficulty in applying Prakken and Sartor's approach to dimensions: that we cannot establish a ranking of pro-plaintiff rules without precedent information. This does not make their representation impossible, or even inappropriate, but it greatly reduces the power of the approach since many of the inferences from a precedent depend on our ability to at least partially order rules with a given tendency. Having objected to Prakken and Sartor, it should be said that neither is it clear how this point is treated in HYPO either; in discussions of on-pointness a quantitative factor such as competitive-advantage is treated as if it were a boolean, simply present or absent, (see for example figure 8.2 on p 132 of Ashley 1990). The quantitative aspects in HYPO do, however, play a very important role in reasoning with the cases when they have been identified: they are used in the generation of hypothetical cases, particularly in order to identify boundary counter-examples, and in distinguishing cases.

It is possible in HYPO to ignore the magnitude of dimensions when considering on-pointness because the quantifiable factors, such as competitive-advantage, can only work in favour of one side. Thus if present competitive-advantage will favour the plaintiff, however weakly. The fiscal domicile case differs in that duration can favour either side, depending on its length, so that the mere presence of duration as a factor cannot be said to favour either side *a priori*. Moreover competitive advantage might also have been such a two-edged factor; it might have been that a very small advantage worked in favour of the defendant rather than very weakly in favour of the plaintiff. Had this been so, magnitude would have need to have been considered in determining on-pointness. In CATO the quantitative aspect of factors is entirely disregarded:

> "HYPO was able to generate arguments comparing cases in terms of the magnitude of their factors ... CATO cannot generate these types of arguments because its case representations do not include the magnitudes of the factors" (Aleven 1997, p211).

What I am suggesting here is that the inclusion of quantitative factors is not a simple extension. In an extreme case it might be that a long enough duration is sufficient to override any other considerations whatsoever, and determine the case for the plaintiff, and a short enough duration is enough to override any other considerations and determine the case for the defendant. But if this is so there are no common sense considerations that can allow us to determine the relative priorities of

```
A & C & Duration = 12 -> p
```

and

```
A & Duration = 20 -> p
```

The point is that the crucial assumption that the more pro-plaintiff factors there are present the stronger the plaintiff's case is violated. Adding C above only strengthens the case if it is of sufficient force to outweigh the weakening consequent on the shorter duration. Of course for a *given duration*, an additional factor will still be sure to strengthen the case. Note that this is true even when the factor favours only one side, as with HYPO's dimensions.

If we cannot establish the rules favouring a particular side without reference to precedents, we cannot form initial partial orders as in Figure 1, and cannot proceed in the manner outlined above. If, as I hope I have suggested above, the plausibility of both Prakken and Sartor and HYPO derived from an appeal to an intuitive sense of such partial orders this is a serious problem.

The problem really stems from the fact that if we have a dimensional factor there is the potential for trading off strength of

satisfaction of such a factor against failure to satisfy others. I shall return to this in section 6, after I have discussed factor hierarchies.

5. FACTOR HIERARCHIES AND SUB-ARGUMENTS

Perhaps the major difference between CATO and HYPO is the introduction of factor hierarchies. This is a recognition that if we consider the factors relevant to a case some seem to belong to one issue and others to another. Factors which address the same area are thus grouped according to their contribution to more abstract factors. For example, in CATO, the abstract factor *Info-Trade-Secret* is determined by the less abstract factors *Efforts-Maintain-Secrecy*, *Info-Known-or-Available*, and *Info-Valuable*. (Aleven 1997, p44.) Note that while the more abstract factor is pro-plaintiff, *Info-Known-or-Available* is pro-defendant. In general an abstract factor is related to less abstract factors favouring both sides. Thus the ascription of the abstract factor typically involves the resolution of a conflict amongst its contributing factors. The use of the factor hierarchy allows CATO to introduce some new argument moves, such as down-playing a distinction where the absence of a contributing factor to an abstract factor is offset by the presence of another contributing factor of the same direction.

Prakken and Sartor model this feature by the use of what they term "intermediate" as opposed to "abstract" factors. If we see the significance of some factors as being the result of abstraction of other factors we need to make a choice between describing a case:

1) using the (relatively few) abstract factors, and reasoning about the ascription of them in sub-arguments using the less abstract factors;

2) excluding the abstract factors (except when making some particular argument moves) and instead including only the most concrete factors, which would be rather like fully unfolding the rules in a logic program;

3) including all factors, concrete or abstract.

(3) does not seem sensible because it leads to double counting; a factor is counted both for itself and for the abstract factor, or factors, to which it contributes. CATO adopts (2), and abstract factors are considered only when producing arguments to evaluate the significance of a particular factor. (1) seems to be favoured by Prakken and Sartor so that the more concrete factors are considered in sub-arguments which once resolved establish whether or not the "intermediate" factor applies to the case.

Is it important which approach we adopt? Consider an example where A and B are pro-plaintiff factors and where D and E are pro-defendant. However, in this case suppose that E can be established on the basis of considering factors E1 and E2. Suppose also that all factors are boolean, and that that E can be established on the basis of *either* E1 or E2 being true. In this case it would seem best to adopt approach (1), since delegation of establishing E to a sub-argument makes things neater and avoids the misleading impression that a case with both E1 and E2 is stronger that one with only one of E1 and E2.

But now suppose that D and E permit degrees of satisfaction (in that they favour their side more if more of their contributing factors are satisfied, rather than being considered only as satisfied or unsatisfied), and that a strongly satisfied E (say with both E1 and E2 true) can compensate for a weakly satisfied D. In such a case we cannot avoid considering both E1 and E2 in conjunction with D, and so (2) seems the best approach.

An interesting point here is that whether E1 and E2 should be considered in a sub-argument or not depends not on anything intrinsic to the factors themselves, but on how the abstract factor which they determine will be used. It is when the strength of satisfaction of E matters that we need to keep the way it was derived under consideration. Contrastingly, if the degree of satisfaction of the abstract feature need not be considered, we should insulate its contributing factors from the rest of the argument, so that we are not misled into over-valuing it. What is particularly interesting is that we can envisage different cases where each of the different treatments would be appropriate. This suggests that there is no single phenomenon involved here.

6. DIFFERENT KINDS OF CONCEPT APPLICATION

We should not be surprised that there are two distinct types of reasoning here. We are familiar with applying concepts for which we can find necessary and sufficient conditions, such as "bachelor", and those for which we must weigh a variety of factors, none of which is decisive, and which permit trade-offs whereby a strongly satisfied factor can counterbalance unsatisfied factors. Such concepts are, since Wittgenstein (1953), often said to exhibit a "family resemblance", and "game" is Wittgenstein's classic example of such a concept (Wittgenstein 1953, section 66).

Legal decisions can also be seen as falling into this division. Consider the concept of pensionable age in UK Social Security Law. Here we have a necessary condition (that the person is over 60 years of age), and two alternative sufficient conditions,(that the person is either over 65, or both over 60 and female). Note that there is no trade off whatsoever between the two factors. Contrast this with a decision as to whether someone is "seeking suitable employment" in relation to the UK Job Seekers Allowance (formerly Unemployment Benefit). In order to determine whether a particular occupation is suitable for a particular individual it is necessary to consider a variety of factors such as the age and state of health of the claimant, the length of unemployment, the nature of the claimant's previous occupation, and the nature of the job on offer. All of these factors must be considered, and some are likely to be favourable to the claimant and some are likely to be unfavourable. None, however, are decisive, and some may be traded off against others.

In a particular decision, both sorts of factors might occur. Suppose that to qualify for Unemployment Benefit a person must:

(a) Be under pensionable age;

(b) Have paid the requisite contributions;

(c) Be seeking suitable employment.

The factors relevant to the decision considered as a whole will be age (relevant to a and c), sex, contributions paid, state of health of the claimant, the length of unemployment, the nature of the claimant's previous occupation, and the nature of the job on offer. One way of looking at the position would be to see it as a family resemblance type decision, with a-c as abstract factors. But another way would be to consider a-c as the factors, and to isolate the family resemblance problems to deciding whether or not c applies. This is possible since there is no trade-off between c and the other two factors. A sixty four year old male who has paid the bare minimum of contributions necessary will qualify if held to be seeking suitable employment, however close the call on this last issue, and a ninety year old will fail however many contributions have been paid and however earnest the search for work. This treatment is also to be preferred since it emphasises the necessity of satisfying the age and contribution conditions.

What this suggests is that when we analyse a case in terms of factors we need to identify not only the factors that contribute to an abstract factor, but also how that abstract factor will be used when required to establish more abstract factors, or reach the final decision. Where the abstract features are used in a straightforward, satisfied or unsatisfied, way they should be established by sub-arguments, whereas if they are used in a family resemblance type decision the contributing factors need to be carried forward into the decision, since it is those that allow us to see how strongly the abstract factor is satisfied.

Now let us return to quantifiable factors. Such a factor could be seen as contributing to a more abstract factor: for example in the fiscal domicile domain the length of duration contributes to the abstract factor *long-duration*. Now if that factor is one where we are only interested in whether it is satisfied rather than the extent to which it is satisfied, we can use the abstract factor, and regard the exact duration as a sub-argument for this factor. If, however, it is to be used in a context in which trade-offs are possible, the extent of the duration must be carried forward and we cannot reduce it to a boolean factor in this way.

7. DISCUSSION

The above suggests that we need to recognise several things when trying to model case-based reasoning using techniques from logical representations of argument.

1) Prakken and Sartor's method of representing precedents does indeed help to account for case based reasoning moves in terms of their argument framework. Difficulties with their approach, centring on factors with quantity and direction which can be involved in trade-offs, are also problematic in HYPO (and ignored in CATO).

2) We need to distinguish decisions in which it is important only whether a factor is satisfied from those in which the degree to which it is satisfied matters as it can be used in trade-offs with other factors. In the first case factors relating to the factor should be considered in a sub-argument, whereas in the second the factor should be unfolded into its components. The importance of this can be seen in CATO. One of its abstract factors is *Information-Trade-Secret*, which would seem to be a *sine qua non* for the plaintiff[4]. If, however, the considerations which lead to this factor are not insulated from the rest of the case, there is nothing to stop them being traded off against defects in the argument for other factors which should be considered independently such as *Confidential-Relationship*.[5]

3) Quantifiable factors can pose problems, especially if they are such that which side they favour depends on their value, since they prevent us from even partially ordering the arguments that can be made for a side *a priori*. This in turn greatly restricts our ability to use precedents to establish definite conclusions as to which side of a case is the stronger. Whether or not they can be eliminated – with respect to deciding the outcome of the case - in favour of more easily handled boolean factors depends on the considerations in (2) and whether it is possible to determine the appropriate cut-off point. Moreover, it may still be undesirable to eliminate quantifiable factors since using only booleans restricts the ability to use arguments (such as slippery slope) which rely on the factor being quantifiable, and which can be seen as attempting to decide where the cut-off point should be. Thus even if we can eliminate quantified factors in favour of booleans, we do so at the cost of losing the opportunity for several kinds of fine-grained case comparison, which are required if we are to reflect actual argument practice. Examples of the importance of arguments based on the quantification of factors in legal argument are given in Rissland (1989).

4) While both methods can help us find the best argument for a given side, and to find objections to this argument, and to reply to those objections, they are not in general able to resolve the issues, but only to identify them. Such systems can therefore clarify the decisions to be made, but not to make them. Nor do they constrain the decisions very much; if a priority is not deducible from the precedents and the partial order, the decision maker has a free choice, although the arguments adduced may be more or less persuasive towards a particular decision.

5) Success of the system remains dependent on a good analysis to identify the factors and to structure them in an appropriate

---

[4] It might, of course, be argued that *Information-Trade-Secret* is not a necessary condition. None the less, necessary conditions do exist, such as being of sufficient age for a retirement pension, and so the need to be able, at least in principle, to insulate factors against trade-off remains, even if CATO avoids this problem in a particular domain.

[5] It was suggested by one of the referees that if one selects cases based on their *overall* similarity (as HYPO does) one is safe regardless of whether the courts trade the factor off or not. I disagree: in the case of retirement pension, two cases may be identical in all respects except that in one the person is 59 and in the other 60. A more similar overall case is difficult to conceive, but the similarities do not affect the decision in any way. The first cannot receive the pension, however favourable the other factors may be, and no precedent with age on the wrong side of the cut-off is relevant.

factor hierarchy. The role of the tools is mainly in making this analysis accessible.

8. CONCLUSION

Throughout its brief history AI and Law has seemed to present a conflict between case based and rule based approaches; panels were held on this at both the first and third International Conferences on AI and Law. Prakken and Sartor's work gives a valuable insight into how these two approaches can be seen as not as in conflict, but as different expressions and ways of implementing the same reasoning phenomena. Differences in flavour result from differences in what is taken as a paradigmatic decision: if we start from statutes we will tend to de-emphasise trade-offs between factors, whereas if we start from cases these will be seen as central. But the truth of the matter is that both are essential, and typical reasoning tasks will involve both kinds of concept application. This needs to be reflected in a system that models this reasoning: any attempt to reduce one kind of decision to the other will lead to anomalies when we attempt to generalise beyond core motivating examples.

9. ACKNOWLEDGEMENTS

I would like to acknowledge some very interesting discussions with colleagues over the years on subjects related to the topic of this paper. In particular I developed some of the ideas in this paper following discussion with Giovanni Sartor, although this is not meant in any way to suggest that he endorses the positions taken here. I would also like to thank the anonymous referees for their detailed comments, some of which I have responded to in this version. I would also like to acknowledge some fruitful Friday conversations with Grant Malcolm, some ideas from which may have found their way in here.

10. REFERENCES

[1] Ashley, K.D., (1990). *Modeling Legal Argument,* MIT Press, Cambridge, Mass.

[2] Aleven, V.A.W.M.M., (1997). *Teaching Case-Based Argumentation Through a Model and Examples*, PhD Thesis, University of Pittsburgh.

[3] Hage, J.C., (1997). *Reasoning with Rules*, Kluwer Academic Publishers, Dordrecht.

[4] Prakken, H., and Sartor, G., (1997). *Reasoning With Precedents in a Dialogue Game*, in Proceedings of the Sixth International Conference on AI and Law, ACM Press New York, pp 1-9.

[5] Prakken, H., and Sartor, G., (1998). *Modelling Reasoning with Precedents in a Formal Dialogue Game*, Artificial Intelligence and Law, Vol 6, Nos 2-4, pp 231-287.

[6] Rissland, E.L., (1989). Dimension-Based Analysis of Hypotheticals from Supreme Court Oral Argument, Second International Conference on AI and Law, ACM Press, pp 111-120.

[7] Wittgenstein, L., (1953). *Philosophical Investigations*, Blackwell, Oxford.