Dynamic Arguments in a Case Law Domain

John Henderson and Trevor 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 we describe an approach to reasoning with cases which takes into account the view that case law evolves through a *series* of decisions. This is in contrast to approaches which take as a starting point a set of decided cases, with no account taken of the order in which they were decided. The model of legal reasoning we follow is based on Levi's account which shows how decided cases often need to be reinterpreted in the light of subsequent decisions, so that features of cases wax and wane in importance. Our aim is to reproduce the arguments that could have been used in a given case, rather than to apply a retrospective understanding of the law to them. A second novel feature is that we use a general purpose ontology to describe the cases, rather than one developed specifically to model the pertinent cases. The paper describes a prototype implementation, and uses an example to illustrate how our approach works. After this case by case description we make some remarks on the insights gained, and draw some conclusions.

1. INTRODUCTION

When a case is decided it is as part of a *process* of development of the law of the domain. There will be relevant past cases, but there will also be future cases. To understand a case properly, it can be argued, we need to see it in this context. This is the approach of writers on legal reasoning such as Levi [5], who explain the reasoning in a particular area of case law through a chronological narrative of a sequence of decisions in that area. The reason for this is that a new decision may cause us to rethink our interpretation of past cases. The same applies to the factors which are used to describe cases in case based reasoning systems such as HYPO [2] and CATO [1]. While at any point in the evolution of the case law we can analyse the decided cases to identify factors, and so provide arguments based on these factors applicable to a new case, we can have no assurance that the decision in the new case will not require us to reinterpret some of the past decisions and the set of factors they should be described under. Moreover it may be artificial to describe the older cases using apparatus that has only emerged much later: certainly, if we do so, we will think of such cases rather differently from those who actually decided them.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ICAIL-2001, St. Louis, Missouri USA.
Copyright 2001 ACM 1-58113-368-5/01/0005 $5.00.

Levi describes the process thus:

> "Reasoning by example shows the decisive role which the common ideas of the society and the distinctions made by experts can have in shaping the law. The movement of common or expert concepts into the law may be followed. The concept is suggested in arguing difference or similarity in a brief, but it wins no approval from the court. The idea achieves standing in society. It is suggested again to a court. The court this time reinterprets the prior case and in so doing adopts the rejected idea. In subsequent cases, the idea is given further definition and is tied to other ideas which have been accepted by courts. It is no longer the idea which was commonly held by society. It becomes modified in subsequent cases. Ideas first rejected but which have gradually won acceptance now push what has become a legal category out of the system or convert it to what may be its opposite. … reasoning by example will operate to change the idea after it has been adopted". [5], pp 5-6.

In this paper we will describe an experiment to model case law in this evolutionary way. Instead of a body of decided cases we will take as out input a stream of cases and their decisions. Instead of factors to describe these cases we will take as a starting point for case description an ontology intended to represent the "common sense" conceptualisation of an area of relevant fact ("the common ideas of society"). As each new case is presented we will produce arguments for an against a position with regard to the case ("the distinctions made by experts"). At any point in the stream of decisions it will be possible to consider a modification of the common sense ontology to represent the current legal conceptualisation. This work represents a development from that reported in [3]. We will use the mechanism described in that paper and use the same example to illustrate the process. In this paper, however, we will introduce a second mechanism and we will focus more on the arguments that are produced in response to each new case, in the context of a prototype implementation.

Section 2 will describe the form of ontology we require for our system, summarise the mechanism described in [3] and suggest a second mechanism for identifying specific relevant cases. Section 3 will introduce the example. Section 4 will describe the prototype implementation. Section 5 will step through the process of applying the mechanisms to generate arguments. Section 6

will offer some discussion of points of interest in this analysis. Section 7 will offer conclusions and prospects for future work.

2. BACKGROUND

2.1 The ontology

Our starting point is a representation of a common sense ontology recording the conceptualisation of the facts of some domain. This is the sort of thing which a product such as Wordnet [6] aspires to provide. A screen shot of Wordnet giving the hyponyms of "professional" (sense 1) is shown in Figure 1.

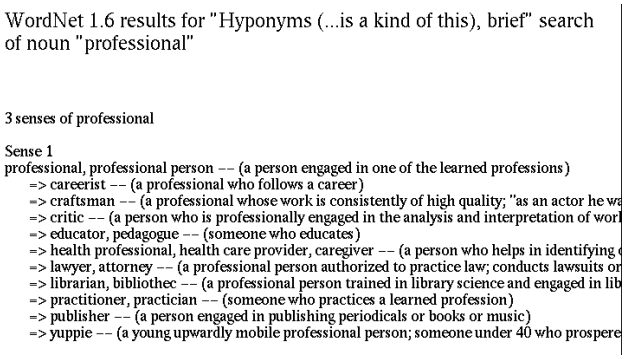
WordNet 1.6 results for "Hyponyms (...is a kind of this), brief" search of noun "professional"

3 senses of professional

Sense 1
professional, professional person -- (a person engaged in one of the learned professions)
=> careerist -- (a professional who follows a career)
=> craftsman -- (a professional whose work is consistently of high quality; "as an actor he wa
=> critic -- (a person who is professionally engaged in the analysis and interpretation of worl
=> educator, pedagogue -- (someone who educates)
=> health professional, health care provider, caregiver -- (a person who helps in identifying
=> lawyer, attorney -- (a professional person authorized to practice law; conducts lawsuits or
=> librarian, bibliothec -- (a professional person trained in library science and engaged in lib
=> practitioner, practician -- (someone who practices a learned profession)
=> publisher -- (a person engaged in publishing periodicals or books or music)
=> yuppie -- (a young upwardly mobile professional person; someone under 40 who prospere

Figure 1: Wordnet hyponyms of "professional"

Essentially this provides an abstraction hierarchy for certain concepts. As it stands, however, is not suitable for our purposes for two reasons. First it rather too inclusive: whilst both yuppies and educators are both professionals, they seem to be professionals in a different sense: to be a yuppie one has to be some other kind of professional, such as a lawyer or an accountant. Second it provides a verbal description which is not easy to compute with. Instead of a verbal description we will supply one or more attributes associated with prototypical instances of that term. An educator for example typically works for some academic institution, has some academic qualifications, a field of expertise, and the like. In our system we use an ontology which is an abstraction hierarchy, pruned to remove extraneous hyponyms, and with each term associated with a set of attributes and values which are held to be true of prototypical members of the class (although *not* all members of the class). As we move down the hierarchy, hyponyms will by default inherit these attributes for their prototypical members, although this inheritance can be cancelled. In our experiments we have used a fairly simple ontology of our own devising. This will be described in section 3. Here we do not use the distinction between distinguishing and prototypical attributes made in [3].

2.2 Finding a general principle

Given such a hierarchy, we can see a case as corresponding to one of the leaf nodes of the hierarchy. The hierarchy can then be used to relate two such cases. If two nodes have a common ancestor, there can be said to be analogies between them; that common ancestor, and all ancestors of that common ancestor, being grounds for the analogies. Two cases may be closely analogous, if their corresponding nodes share a common parent or grandparent, or more distantly analogous. Now suppose that we have a number of decided cases. Any non-leaf node may ground an analogy between two or more cases. Which of these analogies should be considered as significant? What we want is an analogy that gives us predictive power. Suppose we label the leaf nodes representing cases with "p" if they were found for the plaintiff, and "d" if found for the defendant (with respect to the particular issue under consideration). Now any non-leaf node will cover a certain proportion of the decided cases, and of these a certain proportion will be p-cases and the rest d-cases. The ideal analogy for the plaintiff would be one which covered all the p-cases and no d-cases. Typically, however, no such analogy will exist: we will therefore choose the best analogy with respect to *coverage* (the proportion of p-cases covered) and *precision* (the proportion of cases covered which are p-cases). The best analogy will maximise coverage while maintaining precision above a certain threshold. This precision must be greater than 0.5, (so as to make it more likely than not that a case falling under the term will be a p-case), and we have taken 0.66 as our cut off. Once we have found a term to ground our analogy we can consider the d-cases covered by it, and attempt to find a ground for an analogy between these exceptions to the original analogy. This procedure can be applied recursively to find an analogy for the p-cases covered by the exception, and so on. We can finally summarise this as a (defeasible) rule. Suppose the ground of the analogy is A, the ground of the exception B and the ground of the exception to the exception C: the corresponding rule is "p if A and not (B and not C)". A fuller description of this mechanism is given in [3].

2.3 Finding a specific case

The above mechanism gives us an analogy taking into account the whole set of (so far) decided cases. We can also, however, consider the closest analogy to a p-case and a d-case (irrespective of coverage and precision). For a given case this involves finding the closest ancestor which covers (at least) one node corresponding to such a case. Where more than one candidate is found, we choose the node with the most attributes in common. Whereas the mechanism in the above paragraph tries to find a general principle to summarise the body of previous cases, this mechanism finds the best case to cite for a given new case. We will consider both mechanisms when dealing with our examples. We will call the ground generated by the first mechanism the *general grou*nd, and that generated by the second the *specific ground*. Both can be found straightforwardly using standard graph algorithms.

2.4 Generating an argument

The above mechanisms allow us to say things such as "lawyers should be treated like accountants because they are both professionals: since the lawyer is a p-case, so too should the accountant be". This, however, is not necessarily persuasive: what is required is a reason *why* professionals should be p-cases. To generate this we turn to the prototypical attributes. The idea here is that if an applicable term is a good indication of the outcome of a case, it is because things falling under that term have certain characteristics. So the argument is that professionals, for example, should be p-cases because they have certain attributes prototypically true of professionals. We can then counter the argument in a particular case, a barrister for example, by pointing to prototypical attributes of professional which are cancelled in barrister, or by pointing to attributes possessed by prototypical barristers but not by professionals in general. This counter argument can in turn be met by pointing to favourable

cases where the cancelled attribute was also cancelled. The prototypical attributes which are inherited and cancelled can be found simply by following the path from the ground to the case.

The ground of the analogy represents a distinction between cases which can be used to suggest their outcome: the prototypical attributes which emerge from the argument process are important as the basis of factors for describing cases. What we hope will emerge is a set of attributes which are relevant to the decision to be made.

3. AN EXAMPLE

The example is based on the following simplification of reality. Each case in the example concerns the same single question and turns on a single fact. The legal question in the example arises from a series of decisions concerning whether or not a person owes a duty of care to others and are described by the job that the defendant does. For some jobs the duty of care may be owed by the employer of the person giving the advice rather than the individual; in others the nature of the advice given in the job may be such as to give rise to no duty of care at all. For example, under we would suggest that a racing tipster does not owe a duty of care but an investment adviser does. All cases are considered to be of equal authority (ie there are no priority rules between them). The example will be based on 14 cases, in each of which the defendant has a different occupation. Table 1 lists the cases with their outcomes. These outcomes were assigned randomly, to reflect the uncertain nature of legal decisions. We thought that this would offer a fairer test than assigning them according to some preconception of what the law would turn out to be. The question need not be stated as it is the same in every case, namely whether a duty of care was owed.

Table 1: Cases showing Occupations and Outcomes

CaseNo.	Fact	Outcome
C1	Accountant	P
C2	Clerk	D
C3	Solicitor	P
C4	General Practitioner	D
C5	Nurse	D
C6	Lecturer	P
C7	School Caretaker	D
C8	Bank Security Guard	D
C9	School Teacher	P
C10	Builder	P
C11	Broker	P
C12	Homeopath	D
C13	Consultant	P
C14	Barrister	D

These cases need to be seen as the leaves of an abstraction hierarchy. The one we use is shown in Figure 2. Both the set of cases and the abstraction hierarchy are, of course, fictional. We believe that they are, however, sufficiently representative of reality to be of use. Throughout we are interested in modelling a reasoning process, not in providing a legal commentary, and the same process can be applied in a fictitious example as a real one. We have, in fact, also carried out an analysis of actual leading cases relating to duty of care, and so have some grounds for confidence that we are no too much at odds with reality.

We have divided workers into blue collar and white collar workers, according to whether their work is largely manual or not, and also by the field in which they work. We have also distinguished professionals, workers who are valued for their knowledge and expertise, from other workers, whose value lies more in practical skills and their ability to follow more routine procedures. For each term in the hierarchy we now supply some prototypical attributes. These are here intended to be simply illustrative, and to provide a simple example: they make no pretensions to accuracy.

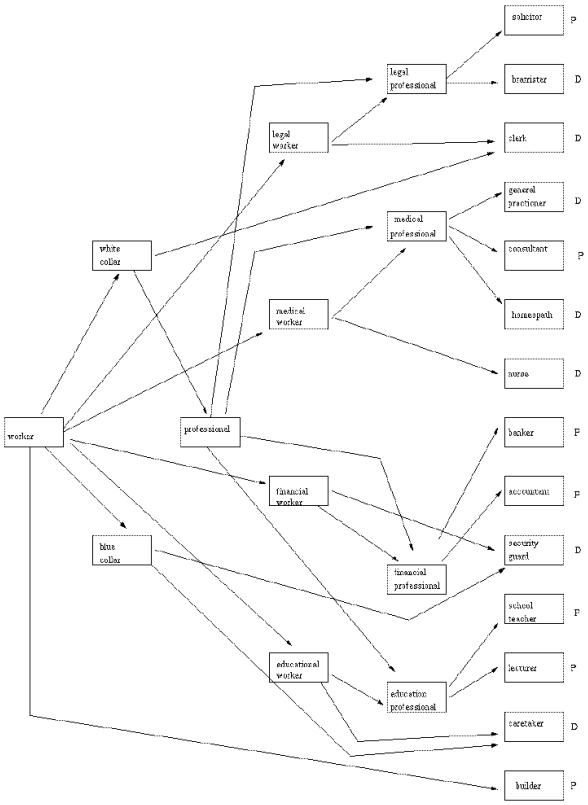

Figure 2: Initial abstraction hierarchy for the cases

Worker is just there to tie the hierarchy together: we will not supply prototypical attributes here. *White collar* and *Blue collar* workers are distinguished by the value of an attribute *works-by*: *hand* for blue collar and *brain* for white collar. *Professionals* are valued for their expertise: typically therefore they have some professional qualification, and often their activities are regulated by a professional association, often the body which issues professional qualifications, (such as the British Medical Association for doctors in the UK). Prototypically also they are employed directly by their client. We give professionals two prototypical attributes: *professional-qualification*, and *employed-by* with the value *client*.

The categories of educational, financial, legal and medical workers are distinguished by the field in which they work. Thus their prototypes will have an attribute *works-in* with values *education*, *finance*, *law*, and *medicine* respectively. Some workers will also be prototypically employed by a particular kind of institution; educators they will have for the values of their employed-by attribute, instances of educational institutions, and medical workers (in the UK) the National Health Service. When we come down to educational and medical professionals we have a case of multiple inheritance which may involve a clash in the values of *employed-by* inherited from the two parents. One of these must be cancelled: for both educational and medical professionals we cancel the *employed-by(client)* inherited from professional. Prototypically, education professionals also have a subject in which they specialise. We include an attribute *specialist-knowledge* to record this.

We now come to the leaf nodes. Barrister cancels *employed-by(client)* and adds *employed-by(solicitor)* as prototypical; consultants and homeopaths cancel *employed-by(national health service)* and add *employed-by(client)*. Consultants have a specialism, and so an additional prototypical *attribute specialist-knowledge* is used to indicate this, while homeopaths cancel the attribute of professional qualification. Lecturers and teachers specialise the employing educational institution to universities and schools respectively. Caretakers and security guards have an attribute *job-description*. Nurses and builders have an appropriate vocational qualification, and builders are also prototypically employed directly by their clients. The situation, for the leaf nodes is summarised in table 2.

In the next section we will describe our current prototype implementation which we have used to explore the above example.

4. Implementation

We have produced a prototype implementation to support and illustrate the above ideas. This comprises a set of tools, designed to produce the information needed to construct the arguments given above. Currently the selection and rejection of the possible lines indicated is done by the user: principles to automate this process will be the subject of future work. The operation will be illustrated with a detailed walk through of the example in the next section.

Throughout we represent the hierarchy as a set of predicates `isa(A,B)`, to be read as "A is a kind of B". Each term in the hierarchy is also associated with a predicate `attributes(Term,[Proto],[Canc])`, where `Proto` is a list of the prototypical attributes associated with that term, and `Canc` is a list of attributes which would otherwise be inherited but are cancelled by the term.

4.1 Finding analogies

The first tool finds the viable analogies, using the notion of the general ground. Given a case, for each superclass of the occupation, we calculate the precision and coverage for plaintiff and defendant analogies, and return for consideration those which meet the threshold of .66 for precision. We then choose, for each side, the analogy which satisfies the precision threshold with the largest coverage. We return this, together with any analogies with greater precision.

Table2: Attributes possessed by leaf nodes. Nodes marked with "*" are introduced by that class rather than inherited

Class	Attribute	Cancelled Attributes
Acc-ountant	Prof Qual works-in(finance) Employed-by(client) works-by(brain)	
Solicitor	Prof Qual works-in(law) Employed-by(client) works-by(brain)	
Barrister	Prof Qual Employed-by(solicitor)* works-in(law) works-by(brain)	Employed-by(client)*
Clerk	works-in(law) works-by(brain)	
GP	Prof Qual works-in(medicine) works-by(brain) Employed-by(NHS)	Employed-by(client)
Consul-tant	Prof Qual works-in(medicine) works-by(brain) special-know* Employed-by(client)*	Employed-by(NHS)*
Homeo-path	works-by(brain) works-in(medicine) Employed-by(client)*	Prof Qual* Employed-by(NHS)*
Nurse	Employed-by(NHS) works-in(medicine) Voc Qual*	
Broker	Prof Qual works-in(finance) Employed-by(client) works-by(brain)	
Bank Security Guard	Employed-by(bank) job description* works-in(finance) works-by(hand)	
Lecturer	Prof Qual works-in(education) Employed-by(univ)* works-by(brain) special-know	Employed-by(client)
School Teacher	Prof Qual works-in(education) Employed-by(school)* works-by(brain) special-know	Employed-by(client)
Care-taker	Employed-by(school)* job description* works-in(education) works-by(hand)	
Builder	Employed-by(client)* Voc Qual*	

The second tool finds the specific analogy each side. The method is to consider the parent of the case: if this has a descendant decided for the desired party, return that descendant: otherwise consider grandparents, and so on, until a suitable case is found. Ties are broken using attributes: the paths from the ancestor to the case and the potential analogy are followed, collecting the prototypical attributes from each node, and where an attribute is cancelled this is moved from this list into a list of cancelled attributes. The lists so collected are then compared.

These two tools form the basis of an argument; we now have, should they exist, both an abstraction of the case and a specific case supporting each side of the dispute.

4.2 Creating Arguments

The third tool attempts to summarise the proposed reasoning as arguments. Two techniques are used, one for the general ground, and one for the specific ground. For the general ground the argument is simply of the form “if ground then plaintiff/defendant”, for each analogy. The opponent can also advance arguments to attempt to show that the particular case is an exception. Where an attribute has been introduced or cancelled below this ground, this is used (unless it is subsequently cancelled or re-introduced): otherwise the class name is used. If there is more than one such attribute, separate arguments are generated for each. These “objecting” arguments are thus of the form “if ground and not exception then plaintiff/defendant".

For the specific ground the argument is of the form “If ground and *attributes in common* then plaintiff/defendant”, where *attributes in common* are shared prototypical attributes introduced below the common ancestor, and negations of prototypical attributes that both have cancelled. If there are no attributes in common, no argument is advanced: the case is considered of no help. Again, objecting arguments can be found: where there are differences between the two cases, these give an argument for the other side of the form "if differences then plaintiff/defendant"

The general ground works with class names, whereas the specific ground makes use of individual attributes. We think there are advantages in this: a class name carries with it connotations, and assumed attributes which may not be explicit in the ontology, but which none the less guide the way its members are thought about. We need, however, also the specific attributes, since we may expect them to capture some of the more important aspects of the cases.

When the case is decided these rules are annotated with the case and the outcome and stored as a predicate:
`arg(Id,[A],C,Prec)` where A is the antecedent, C the conclusion, and Prec is the precedent used in the case of a rule generated from a specific ground, and `gen` otherwise. The `id` as given in this paper is a number in the case of a positive argument, and in the case of an "objecting argument", the number of the argument objected to, an "o" and a number of the objection. (The implementation gives sequential numbers: but we have changed it here in the hope that this will be helpful to the reader.) When we have the decision we also write a predicate which associates the argument id with the case and its outcome. The outcome will be one of "upheld" or "rejected". Note that by "upheld" we mean no more than that the decision is consistent with the argument, and by "rejected" we mean that the argument was defeated by at least one stronger argument. Thus an upheld argument can be rejected in a subsequent case, and a rejected argument accepted, without inconsistency.

These predicates are used by the fourth tool. The fourth tool examines all rules generated by previous cases, and, and for rules in the present case, presents to the user cases in which they were used, and the outcome. Objecting arguments are retrieved only where the initial argument is retrieved. The same rule here is determined only by antecedent and consequent; thus a rule previously used to support the other side may appear. The operation of the tools will become clearer as we consider the example in the next section.

5. The Example in Action

In this section we will go through the example proposed, indicating what our program will produce, and what an intelligent user would do with this information. The narrative will broadly follow that of [3]: where there are significant difference between the arguments given in [3] we will draw attention to them.

When the first case, in our example that of the accountant, is presented, our tools are silent, since there are no precedents to work on. The second case, C2, the clerk, however, allows us to go to work. For the plaintiff there is a single analogy with white-collar worker. There can be no defendant analogies, because there has not yet been a defendant case. Accountant is the specific analogy for the plaintiff, via *white-collar*, and no defendant analogies are possible. C2 therefore, gives the following arguments:
arg(1,[white-collar],p, gen)
arg(1o1,[white-collar,works-in(law)],d, gen)
arg(2,[white-collar],p, C1)
arg(2o1,[not prof-qual,not employed-by(client), works-in(law), not works-in(finance)],d, C1)

Arg1 was generated by the general ground, arg1o1 since clerk may be treated as an exception to this ground because his area of work is added in *legal-worker*. The plaintiff has the specific case C1 which shares only *works-by(brain)*, but there are other attributes not shared with C1 which are put forward in arg2o1. Since C2 was found for the defendant, neither arg1 not arg2 were accepted. We can therefore see *works-in(law)*, and one or more of the exceptions of arg2o1 as potentially able to defeat *works-by(brain)*, and the analogy with white-collar worker. We are next presented with C3, a solicitor. *White-collar* is no longer an analogy since professional has better precision. Indeed, that of white collar is now only 0.5. Professional is, however, an excellent analogy with coverage and precision 1. A solicitor, however, differs from C1 in that he works in law. The defendant has the general analogy of legal-worker. There is a specific analogy for the plaintiff with C1, through *professional*, with the common attributes *professional-qualification* and *employed-by(client)*, and a specific analogy with C2 for the defendant, through *legal-worker*, with common attribute, *works-in(law)*. We therefore get the following arguments.
arg(3,[professional],p,gen)
arg(3o1,[professional, works-in(law)],d,gen)
arg(4,[legal-worker],d, gen)
arg(4o1,[legal-worker,prof-qual],p, gen)
arg(5,[professional],p, C1)
arg(5o1,[not works-in(finance), works-in(law)],d,C1)
arg(6,[legal-worker],d, C2)

arg(6o1,[prof-qual, employed-by(client)],p,C2)

None of these arguments were used in the previous case (5o1 is a weaker version of 2o1). We suggest that the plaintiff would argue that the solicitor is a prototypical professional with a professional qualification and employed by client, as was the accountant. The defendant would rely on the solicitor being a prototypical legal worker, and use the area of work to assimilate the case to C2 rather than C1. In our example the plaintiff wins. The failure of 3o1 suggests that works-in(law) is not a powerful defendant argument, and we may decide accordingly that it was not 1o1 that was decisive in C2. The fourth case is that of a general practitioner. *Professional* is again an analogy for the plaintiff, with coverage 1 and precision 1. The defendant has no general analogy. The plaintiff's specific analogy is with either C1 or C3; we choose C1 as the older case. The defendant's specific analogy is with C2. The case generates the following arguments:
arg(3,[professional],p,gen)
arg(3o2,[professional,works-in(medicine),d,gen)
arg(3o3,[professional,employed-by(nhs),d,gen)
arg(3o4,[professional,not employed-by(client),d,gen)
arg(7,[professional],p,C1)
arg(7o1,[employed-by(nhs), not employed-by(client), works-in(medicine), not works-in(finance)],d,C1)
arg(8,[white-collar], d,C2)
arg(8o1,[prof-qual, employed-by(nhs), works-in(medicine),not works-in(law)],p,C2)

From the previous cases we have arg3 (which was successful in C3). The plaintiff will rely on arg3, but this time the specific analogy of arg7 is not nearly so close as arg5, containing fewer similarities and more objections in arg7o1. The defendants best positive argument is the specific analogy with C2. Note that this is the argument proposed as arg2 by the plaintiff in C2, but since it was rejected there it can now be put forward for the opposing position. There are, however, a large number of differences between C4 and C2, and so the defendant's strongest position is to use arguments 3o2, 3o3 and 3o4 to discredit the analogy in C5. Since the defendant wins C4, we may assume that arg3 is beaten by one of these, or some combination of them. C5 deals with a nurse. The plaintiff this time has no analogy, whereas the defendant has a coverage 0.5 and precision 1 analogy with medical worker. The specific defendant analogy is with general practitioner, whereas the plaintiff has nothing better than C1, which cannot form the basis of an argument because there are no common attributes. The resulting arguments are:
arg(9,[medical-worker],d,gen)
arg(9o1,[works-in(medicine),voc-qual],p,gen)
arg(10,[medical-worker],d,C4)
arg(10o1,[voc-qual,not prof-qual,not works-by(brain)],p,C4)

The plaintiff's case looks weak here: were it to be accepted it would highlight the importance of a professional rather than a vocational qualification. In fact the defendant's case is accepted. so a vocational qualification is not seen as an important exception to medical-worker. C6 deals with a university lecturer. The plaintiff can still use the analogy with professional, (coverage 1 and precision 0.67), and the defendant again has no general analogy. The plaintiff will again use C1 as the specific analogy, but C4, using professional, is a better specific analogy for the defendant than C2. The arguments are:
arg(3,[professional],p,gen)
arg(3o4,[professional,not employed-by(client)],d,gen)
arg(3o5,[professional, works-in(education)],d,gen)
arg(3o6,[professional, employed-by(university)],d,gen)
arg(3o7,[professional,special-know],d,gen)
arg(7,[professional],p,C1)
arg(7o2,[special-know,not employed-by(client),employed-by(university),works-in(education),not works-in(finance)],d,C1)
arg(11,[professional],d,C4)
arg(11o1,[special-know,works-in(education),not works-in(medicine), not employed-by(nhs),employed-by(university),],p,C4)

Again we get arg3, but this time we have three different objections, and, of those which succeeded in C4, only 3o4. The similarities with C1 are exactly those used in C4, and again there are many differences. Note that arg11 relies on these *same* similarities (acceptable because they were rejected in C4, and so it is as yet undecided which side they favour) to urge the defendant's case. Again there are many differences, as indicated by 7o1 for C1 and 11o1 for C4. The plaintiff wins, and so we may think that none of the objections 3o4, 3o5, 3o6 and 3o7 are not good enough to find against professional. This suggests that arg3o4 was not the decisive objection to arg3 in C4. We may also see one of the objections in 11o1 as enough to overturn the analogy with C4.

We now reach a rather different case, C7, the school caretaker. The plaintiff has a general analogy of education worker, with coverage 0.33 and precision 1, and C5 provides the specific analogy, through *education-worker*. The defendant has no general analogy, nor any prior case with any attribute in common. The defendant must therefore rely entirely on objections to the plaintiff's arguments.
arg(12,[education-worker],p,gen)
arg12o1,[education-worker, employed-by(school)],d,gen)
arg(12o2, [education-worker, job-description],d,gen)
arg(12o3. [education-worker, works-by(hand)],d,gen)
arg(13,[education-worker], p,C6)
arg(13o1,[employed-by(school),not employed-by(university), job-description, not prof-qual, not special-know, works-by(hand), not works-by(brain)],d,C6)

In fact the negative arguments prevail; education worker is rejected as the grounds for a plaintiff analogy, and one of the many differences given in 13o1 can be seen as significant. C8 involves a security guard working for a bank. The plaintiff has an analogy grounded in financial worker, with coverage 0.33 and precision 1, and the defendant one grounded on blue-collar worker, with coverage 0.25 and precision 1. The plaintiff's specific analogy is with C1, through financial worker, and the defendant's with C7, through blue-collar worker. We therefore have a number of arguments here:
arg(14,[financial-worker],p,gen)
arg(14o1,[financial-worker, works-by(hand)],d,gen)
arg(14o2,[financial-worker, job-description],d,gen)
arg(14o3,[financial-worker, employed-by(bank)],d,gen)
arg(15,[blue-collar],d,gen)
arg(15o1,[blue,collar,works-in(finance)],p,gen)
arg(15o2,[blue-collar,employed-by(bank)],p,gen)
arg(15o3,[blue-collar,job-description],p,gen)
arg(16,[financial-worker], p,C1]
arg(16o1,[employed-by(bank),not employed-by(client), works-by(hand),not works-by(brain),job description, not prof-qual],d,C1)
arg(17,[blue-collar, job-description],d,C7)

arg(17o1,[works-in(finance), not works-in education), employed-by(bank),not employed-by(school)],p,C7)

Here again the defendant wins. In the last two cases, the arguments for the defendant generated deviated from those that we suggested would be best in [3]. There we suggested that the simple argument "not a professional" would suffice. That such an argument is not generated is because the ontology does not have such a class. If such a class did exist, it would ground a defendant analogy with coverage 0.75 and precision 1, and so would have been reported. The lack of such a class reflects a presumed uselessness of the distinction in common life: however, that it would give rise to the natural legal argument suggests that a rather different legal conceptualisation is forming as the cases develop. We will return to this point later. C9 is the school teacher. The plaintiff has analogies with professional (coverage 1 and precision 0.75), and educational professional (coverage 0.33 and precision 1). The defendant has no analogies (educational worker since C7 having a precision of 0.5 for both sides). The specific analogies are with C6 for the plaintiff and C7 for the defendant. The arguments are therefore:
arg(3,[professional],p,gen)
arg(3o4,[professional,not employed-by(client)],d,gen)
arg(3o5,[professional, works-in(education)],d,gen)
arg(3o7,[professional,special-know],d,gen)
arg(3o8,[professional, employed-by(school)],d,gen)
arg(19,[education-prof],p,gen)
arg(19o1,[education-prof,employed-by(school)],d,gen)
arg(20,[edu-prof],p,C6)
arg(20o1,[employed-by(school),not employed-by(university), d,C6)
arg(21,[education-worker,employed-by(school)],d,C7)
arg(21o1,[prof-qual, not job-description, works-by(brain, not works-by(hand),special-know)],p,C7)

In C6, none of 3o4, 3o5 and 3o7 prevailed over 3, and so the defendant must urge 3o8 as the main objection, which is also the basis of the objection to arg19. This is also a key attribute in common with C7, the defendant's specific case. None the less, considering arguments 20 and 21 and their objections, it is not unreasonable to see C6 as closer than C7, and we may not be surprised when the plaintiff wins. Had the result gone the other way, *employed-by(school*) would have assumed considerable importance.Next we get a very different case, C10, the builder. Builder has only worker in common with any prior case, and so there are no analogies with sufficient precision. For the specific analogies, all cases are equally close and so we must look at attributes. The builder has only one attribute in common with a prior defendant case, the vocational qualification shared with C5. He also has employed-by(client) in common with the plaintiff cases C1 and C3. We choose C1 as the earlier case.
arg(22,[worker,employed-by(client)],p,C1)
arg(22o1,[voc-qual, not prof-qual, not works-in(finance),not works-by(brain)],d,C1)
arg(23,[worker,voc-qual],d,C5)
arg(23o1,[employed-by(client),not employed-by(nhs),not works-in(medicine),p,C5).

The case is a hard one, because the best established analogy, professional, has two prototypical attributes, and whereas in the past we had not met employed-by(client) without prof-qual, here we have such a case. When this is decided for the plaintiff, it seems to suggest that employed-by(client) is crucial, as suggested by arg22. C11, the broker, is by contrast straightforward. The plaintiff can use analogies with professional (coverage 0.8 and precision 0.84) and with financial professional (coverage only 0.2, but precision 1). The defendant has no general analogy. The specific analogy for the plaintiff is with C1, and for the defendant with C8. The arguments are:
arg(3,[professional], p,gen)
arg(3o9,[professional, broker],d,gen)
arg(24,[financial-prof],p,gen)
arg(24o1,[financial-prof,broker,d,gen)
arg(25,[financial-prof],p,C1)
arg(26,[financial-worker],d,C8)
arg(26o1,[prof-qual,not job-description, employed-by(client), not employed-by(bank), works-by(brain), not works-by(hand)],p,C8)

This looks rather straightforward. All the defendant can argue is that like C8, C11 works in finance, but this is a similarity with the plaintiff case C1, and was used as a pro-plaintiff argument in C8. The only objections to the general analogy are that there is something special about brokers, not represented in the ontology. In our example the obvious happens, and the plaintiff wins. In C12 we have an example of the homeopath, who has no professional qualification, but is employed by the client. The analogy with professional works for the plaintiff (coverage and precision 0.84), but the defendant has the analogy with medical worker, with coverage 0.4 and precision 1. The defendant has a specific analogy with general practitioner. The plaintiff's specific analogy is interesting: our algorithm gives C1, although a case might be made for C10, builder, which also lacks the professional qualification. Again this is partly a consequence of our representation where the semantic connection between vocational qualification and lack of professional qualification cannot be made. The arguments here are:
arg(3,[professional], p,gen)
arg(3o10,[professional,not prof-qual],d,gen)
arg(3o2,[professional,works-in(medicine),d,gen)
arg(27,[medical-worker],d,gen)
arg(27o1,[medical-worker, employed-by(client)],p,gen)
arg(27o2,[medical-worker, not employed-by(nhs)],p,gen)
arg(27o3,[medical-worker, works-by(brain),p,gen)
arg(28,[professional],p,C1)
arg(28o1,[not prof-qual, works-in(medicine),not works-in(finance], d C1)
arg(29,[medical-prof],d,C4)
arg(29o1,[not pro-qual, employed-by(client),not employed-by(nhs)],p,C4)

Here, although after C10 we might think that 27o1 was a good enough objection to arg27, the case is found for the defendant. This means that we need to give respect to 3o10 and 3o2. Recalling that 3o2 was upheld in C4, we may begin to think that working in C13 offers another kind of medical worker, a consultant. the plaintiff has no better analogy that professional, with coverage still 0.84, but precision now down to 0.71. The defendant again had medical worker, coverage now up to 0.5, and precision 1. The specific analogy for the defendant is again C4, but the plaintiff can this time use C6 which has more attributes in common than C1, because of to the specialist knowledge of the consultant. This gives the following arguments:
arg(3,[professional], p,gen)
arg(3o2,[professional,works-in(medicine),d,gen)
arg(3o7,[professional,special-know],d,gen)
arg(27,[medical-worker],d,gen)
arg(27o1,[medical-worker, employed-by(client)],p,gen)
arg(27o2,[medical-worker, not employed-by(nhs)],p,gen)

arg(27o4,[medical-prof,special-know],p,gen)
arg(30,[professionall, special-know],p,C6)
arg(30o1,[employed-by(client), not employed-by(university), works-in(medicine) not works-in(education)],d,C6)
arg(31,[medical-prof],d,C4)
arg(31o1,[special-know, employed-by(client), not employed-by(nhs),p,C6]

Since this uses 3 and 27 as did C12, we can make some instructive comparisons since C13 is decided for the plaintiff. 3o7 was rejected in C6 and is rejected here also. But 27o4, is new to C13, and this seems enough to overturn the previously favoured analogy in 27, drawing attention to the importance of *specialist-knowledge*. Note also that 3o2, successful in C4 and C12, fails here. The last case, C14, concerns a barrister. The general plaintiff analogy is still professional, now with coverage 0.86 and precision 0.75, and also legal professional with precision 1. The defendant has no general analogy. The plaintiffs specific analogy is with the legal professional, C3, whereas the defendant can go to either C2, through legal worker, or to C4 or C12 through professional. On attributes C4 is closest. The arguments now are:
arg(3,[professional], p,gen)
arg(3o1,[professional, works-in(law)],d,gen)
arg(3o4,[professional,not employed-by(client)]d,gen)
arg(3o11,[professional, employed-by(solicitor)],d,gen]
arg(32,[legal-prof],p,gen)
arg(32o1,[legal-prof, not employed-by(client)]d,gen)
arg(32o2,[legal-prof, employed-by(solicitor)],d,gen)
arg(33,[legal-prof)],p,C3)
arg(33o1,[employed-by(solicitor), not employed-by(client)],d,C3)
arg(34,[professional,not employed-by(client)],d,C4)
arg(34o1,[works-in(law), not works-in(medicine), employed-by(solicitor),not employed-by(nhs)],p,C4)

This final case is found for the defendant, suggesting that 3o11 might be important, since 3o1 failed in C3 and 3o4 in C6 and C9.

6. Discussion

We will now make some brief remarks about what we have learned from this example.

6.1 Reinterpretation of cases

One of the important aspects of seeing case law as a process was that we could reinterpret cases in the light of subsequent decisions. Let us see the extent to which this was borne out.

When C1 is decided many interpretations of its significance are possible: the accountant could be representative of white collar, financial worker, professional, or simply in a class of its own. The facts of the next case require the broadest interpretation, white collar, to be made by the plaintiff. This is rejected in C2, however, thus narrowing the possibilities for interpreting C1. C2 also suggests a number of features that might have significance for the defendant: the lack of a professional qualification, not being employed by the client, working in law and not working in finance. Again we cannot yet say which is crucial. In C3, C1 is interpreted as an example of professional. But C3 dismisses the objection based on working in law. We can now revisit C2, and see the attribute *works-in(law)* as not significant. This tells further against the significance of *white collar* in C1, suggesting that *professional* is a more appropriate abstraction for C1. The differences between C2 and C1, since the area of work is downplayed. suggests that the important attributes are *professional qualification* and *employed-by(client).* Since these are the prototypical attributes of a professional, everything seems to hang together. This interpretation can accommodate C4, which, although involving a professional was found for the defendant, cancels one of the prototypical attributes. C5 also fits, lacking both the central attributes.

C6, however, which cancels the same attribute as C4, is found for the plaintiff, and so we could reinterpret C4 and C5. Under a new interpretation we could see the professional qualification as the important aspect of professional, and the area of work or the specific employer, the NHS, as being the reason why professional is not enough in C4 and C5. C7, C8 and C9 fit nicely with this interpretation, and the former two suggest *blue-collar* as important for the defendant. If this means that *works-by(hand)* is sufficient for the defendant, *works-by(brain),* which has appeared insignificant since C2, now assumes some importance as indicating that *works-by(hand)* does not hold. C10, however, forces further reconsideration. It appears that *employed-by(client)* is decisive in finding for the plaintiff in C10. So this can again be seen as the reason why C4 was decided as it was, and the area of work and the employer as of lesser importance. But now we need to review C6 and C9 to explain why those cases, without *employed-by(client),* were decided for the plaintiff. From the arguments in these cases it seems that employment in education or the specialist knowledge must be the reason to exempt the education professionals. But these arguments were in fact urged by the defendant, and only used by the plaintiff to distinguish them from C4. Moreover the former failed in C7. We might therefore now recast C6 and C9 as being crucially dependant on specialist-knowledge for their decision for the plaintiff, even though *at the time*, this attribute was used to argue for the defendant. C11 adds little but confirmation of our understanding. But C12 is a case with *employed-by(client),* but which is found for the defendant. The most plausible reason for this seems to be the lack of a professional qualification. Since *employed-by(client)* is no longer a sufficient condition for the plaintiff, we may now want to look at C10 again, and see a vocational qualification as an acceptable alternative to a professional qualification. But, if this is so, why were C4 and C5 found for the defendant? In the light of C6 and C9, we can explain this by their lack of specialist knowledge, or by their area of work, or by their specific employer. C13, which has all three of the pro-plaintiff attributes including specialist knowledge seems to support the first interpretation, rule out the second, and is neutral as to the third. But now C14, which follows C4 and is found for the defendant, has a different employer, while lacking specialist knowledge, which seems to argue for specialist knowledge as what matters. We thus arrive finally at an interpretation which covers all cases:

> *GP1: Plaintiff if and only if professional or vocational qualification and (employed-by(client) or specialist-knowledge)*

This explains all the cases, but was only apparent when C14 was decided, and could not be argued before C13, since *specialist-knowledge* became important only when C6 and C9 were re-interpreted in the light of C12. Moreover, we can have no assurance that this principle will not need to be revised when subsequent cases are considered.

All in all the example as modelled provides good evidence to support the case for re-interpreting cases in the light of subsequent decisions. An important test given a set real cases would be to examine the texts of the decisions to see whether the

attributes thrown up by our model were in fact used in the reasoning of the judges.

6.2 The Importance of Case Order

A second point we wanted to make was that the order in which cases were decided was important for the way in which they were considered. Space precludes detail on this point, but for illustration suppose we decided the cases in the reverse order, and consider the arguments in C9.

Here we have no general analogies for either side, precision on professional being 0.5. In passing we can note that professional which emerged early and strong as the central feature of plaintiff arguments from the earlier sequence here will have appeared in C13 and C11 - but as an argument for the defendant. The specific plaintiff analogy is now with consultant, and the specific defendant analogy is with barrister. The arguments therefore are:
arg(35,[professional,special-know], p,C13)
arg(35o1,[works-in(education),not works-in(medicine), employed-by(school),not employed-by(client)],d,C13)
arg(36,[professional,not employed-by(client)],d,C14)
arg(36o1,[employed-by(school),not employed-by(solicitor), works-in(education), not works-in(law)],p,C14)

These are very different from the arguments advanced for C9 as part of the original sequence. One important effect of deciding this for the plaintiff is to recognise the importance of specialist knowledge for the plaintiff at a much earlier stage. The arguments from this new sequence remain rather different throughout. Professional, now will become available for the plaintiff only when considering C4 - when it will be rejected, and only become available again *after* C1, and so will not reappear in this sequence. Since this was a central consideration in the previous sequence, this in itself draws attention to the importance of the order in which the cases are considered.

6.3 General versus Legal Specific Ontologies

If we were to model an ontology after all the categories had been decided, when we had formed GP1 as our overall interpretation, we would choose very different categories. We would want categories to group those with qualifications (whether professional or vocational), those employed by the client as against any other employer, and those with specialist knowledge, whether medical or educational. Such an ontology would produce better analogies, and require fewer exceptions. We might even produce a class which prototypically had all three attributes and call it "professional in the legal sense". It is certainly true that as case law evolves the conceptualisation used by the law - which of necessity started from the common sense conceptualisation - develops and becomes increasingly precise. Distinctions, not necessarily useful in everyday life, become the points on which cases turn.

At some point, in may be useful, in arguing a case, to put forward a new conceptualisation explicitly, since it will make emergent distinctions sharp. An interesting problem for future investigation would be to try to automate the construction of such a revised ontology as cases come in. We would, however, defend our current use of a static common sense ontology both since it is the only available starting point, and since any revised ontology would itself be open to revision. Our current view is that such a revised ontology is more useful to clarify the understanding of the reasoning of a body of decided cases than in arguing current cases.

6.3 Relation to Previous Work on CBR

Previous work on Case Based Reasoning can be broadly considered as using one of two basic approaches: the feature based approached where the features of a case are considered and the closest case selected, such as LASER [7] and a factor based approach, exemplified by HYPO [2] and CATO [1], where cases are considered according to issues which are held to favour either the plaintiff or the defendant. Both approaches work from a body of cases. We will make some remarks on both.

The first approach will pick out the same cases as the specific analogy - provided the analogous case was decided before the case being considered. One result of this is that a case may retrospectively seem simple, because it is close to cases that were decided subsequently, using it as a precedent, whereas when it was decided because these similar cases were not available, it seemed to be a hard case. Such a system can give good results where the law is fairly stable, and there are sufficient cases decided to provide satisfactory precedents for cases which arise. We may, however, doubt that this is ever the case: what matters is not only whether the new case differs from prior cases and whether these differences *have been* found to be significant, but also whether they *will* be found significant. In our example specialist knowledge was found to be significant, but only some time after the cases which relied on it had been decided. If we had considered the law settled after C11, we might well discount this factor when matching, and so be misled.

Turning to HYPO like systems, we want to make several remarks. First, the factors and dimensions used to describe the case will be determined in accordance with the understanding of the analyst at the end of the sequence. For the example, *qualified* and *specialist-knowledge* might be used as pro-plaintiff factors and *not-employed-by-client* and perhaps *employed-by(nhs)* as pro-defendant factors. These factors will generate acceptable 3-ply arguments from the set of cases. For example consider C14. The plaintiff will cite C3, and the defendant can respond by citing C4 as a counter-example, and distinguishing C14 from C3 since C14 is not employed by client. The plaintiff will distinguish C4 on employed-by(nhs), and cite C6 and C9 to show that the weakness of the employer is not fatal. But because understanding evolves, the set of factors that are chosen depends on when we do the analysis: after another dozen cases, a different set of factors may have emerged. The system thus relies on the analysis being done at an appropriate time, when a good set of factors can be found. Second, we want to say something about abstract factors, as used in CATO. These appear to be of two different kinds. One might be an abstraction taken from the common sense ontology: thus "professional" could be used as an abstract pro-plaintiff factor. Other abstract factors relate to the ontology as conceived specific to this area of law, and would include factors such as "qualified". The third point is that because factors may emerge subsequent to a case being decided, many decisions will not be written in terms that can be related to the factors. This is a problem is we want to detect factors relevant to a case from the texts of the decisions. Because we are imposing our analysis retrospectively, we cannot expect the decisions to be framed in terms of that analysis.

Mention should also be made of [9], which attempts to identify concept drift in a stream of cases, by identifying structural changes in the decision trees induced by a machine learning

algorithm. Here cases are represented as a pre-defined set of case facts, and no use of a common-sense ontology is made. A more detailed comparison of this work may be instructive as future work.

7. CONCLUSIONS

In this paper we have describe a prototype implementation of a model of legal reasoning, which views this activity not as the application of a settled body of law derivable from a set of cases, but as a process in which cases influence future decisions, but are themselves clarified and reinterpreted in the light of subsequent decisions. Moreover we base our arguments not on a purpose built conceptualisation of the domain, but on an ordinary, general purpose, common sense conceptualisation. Advantages of our approach include:

- There is no need of a special legal analysis of the domain: instead we use a general conceptualisation. Current work, such as Wordnet [6] and CYC [4], suggest that at some future date such a conceptualisation will be pre-existing and available for re-use. If it does need to built from scratch, no special legal expertise is required.
- We believe we have a realistic model of the reasoning process, both in that cases use only materials available when they were considered, and because the reasoning can be tied to common concepts rather than to specially produced terms of art.
- The maintenance problem is avoided: since we do not use any particular set of cases to ground an analysis, there is no possibility of that analysis being overturned by subsequent decisions.
- The system can produce arguments, rather than simple case citations. These arguments are natural, because derived from a common sense conceptualisation.
- Our computational mechanisms are very simple and straightforward.

There are, of course, some disadvantages, as well. These include:

- The arguments produced may or may not be sensible. Attributes we might take to favour a particular side of the dispute can appear in arguments for the other side. In consequence some manual filtering of the arguments is required.
- Because we start from a general conceptualisation, some concepts that will emerge as important, may not appear in the ontology, which can give a rather artificial feel to some of the arguments. This suggests that at some point in the process, we might need to think about transforming the ontology to capture concepts peculiar to the law of the domain.
- Because the system generates all possible arguments, the richer the ontology the more arguments appear. Many of these will not be sensible. Even in our simplified example, some attributes, such as the area of work, are probably unnecessary noise, and in any putative real ontology, there are likely to be more attributes, exacerbating this problem. Possibly the ontology will need some prior analysis to identify potentially relevant attributes, to reduce the number of spurious attributes.
- The ontology lacks semantics for the attributes. This can hide some useful distinctions and connections such as that between employed-by(client) and having another employer, or the potential connection between vocational and professional qualifications. Also in general the complement of the extension of an attribute may be as useful as the extension itself, but this is hard to pick up from the kind of ontology envisaged.

Even given these disadvantages, we believe that the approach taken here is interesting, and offers some insights not available in existing approaches which ignore the process aspects of reasoning with cases. We intend to investigate further. Possible areas of future work include:

- Refining our implementation. In its current state it is neither robust, nor well integrated. We would hope to improve this situation, and package it up with an acceptable user interface;
- Applying some rudimentary techniques to present the arguments in a more readable form, rather than as a list of predicates;
- Trying to identify some heuristics for identifying bad arguments. This may well require some additional information, describing any expectations about which party an attribute will favour;
- Exploring what use can be made of the absence of attributes, as well as positive occurrences. This may require some adaptation to the attributes used in the initial ontology;
- Exploring the possibility of automatically producing a refined ontology with categories which emerge as interesting during the process of deciding cases, to produce interim and revisable conceptualisations.

10. REFERENCES

[1] Aleven V., *Teaching Case Based Argumentation Through an Example and Models*. PhD Thesis. The University of Pittsburg. 1997.

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

[3] Henderson, J., and Bench-Capon. T., (2000). A model of the the development of distinctions in case law, in Beuker, J., Leenes, R., and Winkels, R., *Legal Knowledge and Information Systems, Jurix 2000*, Amsterdam, IOS Press. pp23-34..

[4] Lenat, D.B., and Guha, R. V., (1990). *Building Large Knowledge Based Systems: Representation and Inference in the CYC Project*: Addison Wesley

[5] Levi, E.H., (1948). *An Introduction to Legal Reasoning*, University of Chicago Press.

[6] Miller, G.A., Beckwith, R., Fellbaum, C., Gross, D., and Miller, K., (1993). *Introduction to WordNet: An On-Line Lexical Database*: Cognitive Science Laboratory, Princeton University

[7] Montazeri, M.A., Bench-Capon, T.J.M., and Adam, A.E., (1997). LASER: a System to Retrieve UK Employment Law Cases. *Informations and Communications Technology Law*, Vol 6 No 1, pp 41-54.

[8] Prakken, H., and Sartor, G. (1996). A Dialectical Model of Assessing Conflicting Arguments in Legal Reasoning. *Artificial Intelligence and Law* 4: 331-368.

[9] Rissland, E.L., and Friedman, M.T., (1995). Detecting Change in Legal Concepts. In Proceedings of the Fifth International Conference on AI and Law. ACM Press: New York, 127-136.