Artificial Intelligence and Law manuscript No.
(will be inserted by the editor)

Accommodating Change

Latifa Al-Abdulkarim · Katie Atkinson · Trevor Bench-Capon

the date of receipt and acceptance should be inserted later

Abstract The third of Berman and Hafner's early nineties papers on reasoning with legal cases concerned temporal context, in particular the evolution of case law doctrine over time in response to new cases and against a changing background of social values and purposes. In this paper we consider the ways in which changes in case law doctrine can be accommodated in a recently proposed methodology for encapsulating case law theories (the ANGELIC methodology based on Abstract Dialectical Frameworks), and relate these changes the sources of change identified by Berman and Hafner.

Keywords Legal Reasoning, Factors, ADF, Knowledge Engineering, Argumentation, Values

1 Introduction

Although Carole Hafner worked on several topics over her long career, first on conceptual retrieval [33], and, later, on ontologies [40] and legal drafting [35], many believe that the chief jewels in her crown were the series of three papers on reasoning with legal cases that she wrote with Don Berman for ICAIL between 1991 and 1995: [20], [21] and [22], consolidated in Carole's paper for the issue of *AI and Law* in memory of Don [34]. In the first half of the nineties, reasoning with legal cases was the main focus of AI and Law research in the US (in Europe, systems using rules based on formalisations of legislation, such as [48], were more in vogue at the time) and HYPO [7], CATO [5] and CABARET [49] were the leading case based systems. In all three of these systems the cases in the case base were considered as being homogeneous and equal in status despite being drawn from *different times* (from 1945 onwards in CATO), *jurisdictions* (a variety of states in CATO as well as the Supreme Court) and *levels of court*, from circuit courts to the Supreme Court.

Latifa Al-Abdulkarim
Department of Computer Science. University of Liverpool. E-mail: latifak@liverpool.ac.uk

Katie Atkinson
Department of Computer Science. University of Liverpool. E-mail: katie@liverpool.ac.uk

Trevor Bench-Capon
Department of Computer Science. University of Liverpool. E-mail: tbc@liverpool.ac.uk

While treating the cases as homogeneous is one valid way of doing research (philosophers often discuss arguments of their long dead predecessors as if they were contemporaries in the SCR), and allows the collection of a reasonable number of interesting cases, the three papers of Don and Carole explored how ignoring the distinctions of time, place and procedural context could distort the ways in which precedents were considered, and the roles played by them in new decisions.

The 1991 paper [20] examined the effect of taking procedural context seriously. This aspect has not been much explored since, although such information is sometimes included in legal ontologies such as [52], and the modelling of different levels of court in argument frameworks was explored in [50].

The second paper [21], which introduced the notion of the purposes of the laws, or the social values they were designed to promote, as a way of guiding decisions in cases not yet covered by precedents, has proved highly influential, giving rise to a whole series of papers exploring the use of social values in legal case based reasoning. These include several papers from the issue of *Artificial Intelligence and Law* in memory of Don Berman [19], [41] and [47], later work such as [18], and very recent work such as [42], [6] and [1].

The third of the Berman and Hafner papers, [22], looked at how changes over time could affect case law, and in particular identified indicators that the case law was moving out of a period of stability [38] and that precedents were likely to be refined or even overturned (called *red flags* in [22]). This was also the topic of [45], also from ICAIL 1995, and the later [46]. In this paper we will focus not on detecting that change is coming but on accommodating change in an existing representation of case law. The representation we discuss will be that produced by the ANGELIC methodology [2], a recent proposal based on Abstract Dialectical Frameworks (ADF), a generalisation of Dung's argumentation frameworks [32] introduced in [26] and refined in [25]. That changes present problems for systems based on statute law had already been recognised in [24], and that systems needed to be designed with change in mind, was the motivation behind so-called isomorphic representations [14], [15].

1.1 ADFs for Case Law

For our exploration we will suppose that we have have a body of case law represented as an Abstract Dialectical Framework (ADF) [25], as described in [2], which offers a full description of how ADFs can be used to encapsulate a body of case law (the *ANGELIC* methodology). Formally an ADF can be defined [25] as:

Definition 1: An ADF is a tuple $ADF = < S, L, C >$
where S is the set of statements (positions, nodes), L is a subset of $S \times S$, a set of links:
and $C = \{C_{s \in S}\}$ is a set of total functions $C_s : 2^{par(s)} \rightarrow \{t, f\}$, one for each statement s.
C_s is called the acceptance condition of s.

The ANGELIC methodology partitions L into links supporting the parent ($L+$) and those attacking the parent ($L-$). As an example of an ADF encapsulating case law we will use the Wild Animals cases introduced in [21] and subsequently extended to include the famous *Popov v Hayashi* case [51], and further refined in [12]. The ADF is shown in tabular form in Table 1, and in diagrammatic form in Figure 1.

The acceptance conditions for the nine non-leaf nodes are:

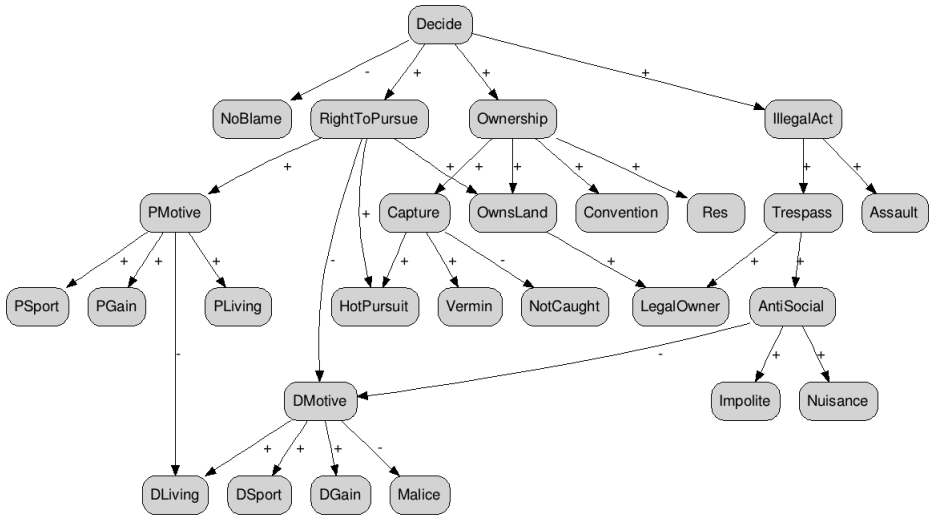

Fig. 1 ADF Diagram for Wild Animals

S	L+	L-
Decide for Plaintiff	Ownership,RightToPursue,IllegalAct	NoBlame
Capture	HotPursuit, Vermin	NotCaught
Ownership	Convention, Capture, OwnsLand, Res	
PMotive	Pliving, PSport, PGain	DLiving
DMotive	DLiving, DSport, DGain	Malice
OwnsLand	LegalOwner	
RightToPursue	OwnsLand, Pmotive, HotPursuit	DMotive
AntiSocial	Nuisance, Impolite	DMotive
Trespass	LegalOwner, AntiSocial	
IllegalAct	Assault, Trespass	

Table 1 Wild Animals as ADF

```
Decide for Plaintiff if NOT (NoBlame)
        AND ((Ownership
            OR (RightToPursue AND IllegalAct)).
Ownership if (OwnsLand AND Resident)
        OR Convention OR Capture.
Capture if NOT (NotCaught) OR (Vermin and HotPursuit).
RightToPursue if OwnsLand OR
           ((HotPursuit AND PMotive
           AND (NOT (better) DMotive)).
PMotive if PLiving OR
           ((PSport OR PGain)
           AND (NOT DLiving).)
DMotive if NOT Malice AND
           (DLiving OR DSport OR DGain).
IllegalAct if Trespass OR Assault.
Trespass if LegalOwner AND AntiSocial.
AntiSocial if (Nuisance OR Impolite)
```

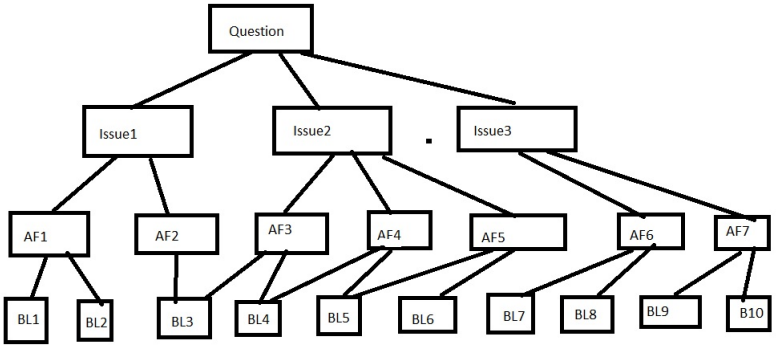

Fig. 2 Schematic Case Law ADF

```
AND (NOT DLiving).
```

Informally ADFs have the advantages of combining a hierarchical structure of concepts, as found in the abstract factor hierarchy of CATO [5], with a traditional rule based representation of the acceptance conditions as found in work going back to [48]. Moreover ADFs provide the rule base with a useful and sensible partitioning [1], which makes for an effective modularisation of the design. The software engineering benefits of using ADFs as a design tool for constructing sytems based on case law are further discussed in [2].

As well as the Wild Animals domain, two other examples of ADFs encapsulating knowledge of case law are given in [2], relating to US Trade Secrets as modelled in CATO [5] and IBP [27] and the automobile exception to the US Fourth Amendment, used in [44] and [9]. In this paper, however, since we are interested in structural modification of the ADF, we will use the purely abstract example shown in Figure 2. Like the examples in [2], and following the principles of [4], there is an overall question tying the structure together, and layers for representation of issues, abstract factors and base level factors.

2 Types of Changes

In [45], which represents case knowledge as a binary decision tree, a number of different types of changes to the law are described in terms of changes to the structure used. A concept could be *generalised* (either by adding a disjunct or removing a conjunct) or *restricted* (either by removing a disjunct or adding a conjunct). The *relevance* of an attribute can change by moving up (or down) the tree. Finally the value of an attribute can be *inverted* by shifting it from the true branch to the false branch. Looking at changes in terms of modifications to the structure of the underlying tree is an excellent way of systematizing our investigation.

Returning to [22], which represents cases in terms of *dimensions*, following HYPO [7], the solution proposed was to associate *purposes* (as introduced to legal CBR in [21]) with the various dimensions. Note that this differs from the way values were used in much of the later work by other people which built on [21], such as [18] where values were associated with factors. Thus an important part of the background knowledge from which theories were constructed in [18] was a set of factors, each associated with a value and the party the presence of the factor favoured. For individual cases, the facts determine a specific point on the dimensions, which in turn determines which party to the dispute is favoured on each

dimension. This is very similar to the proposals in [42] and [3] in which points on dimensions are used to bridge between facts and factors, and [16], which uses facts represented as dimension points to link to values. Seeing changes as resulting from additional or changing purposes and value preferences was the key to the solution proposed in [22].

The final addition to the representation of cases made in [22] is that the outcome is not just the final result, but also a set of *holdings*. These intermediate conclusions can be seen as issues, or as abstract factors, and recognise that the change will affect particular aspects of the cases. This requires the ability to reason with portions of precedents in the manner of Branting [23], and the holdings provide a way of dividing cases into sensible portions. As discussed in [1], the ADF also localises the effect of precedents, and considers only a relevant subset of factors from precedents, so that there too the focus is on portions of precedents, rather than the whole cases. The holdings of [22] can be seen as corresponding to the abstract factor and issue nodes of the ADF, and supplying their acceptance conditions.

Thus we find that [22] identifies several additional components which should be included in the case representations to enable the modelling of changes in case law. The additions were ahead of their time: it took more than five years before the importance of intermediate concepts was properly recognised [8], [10], [39], and even longer for the role of dimensions in bridging from facts to factors [42], [13], [3] and [4] to be seen as essential.

3 Amending an ADF

In [2] modification of the ADF was discussed in the context of refining a program to provide better performance against a set of test data. These modifications involved either altering the priorities in the acceptance tests, or adding a node (supporting or attacking). In that paper we did not consider the removal of nodes, on the basis that the same effect can be achieved by lowering the relative priorities, but for changes in the law, whether statute law or case law doctrine, removal is an additional possibility. A key difference between the ADF of [2] and the decision tree of [45] is that the children of a node in an ADF do not distinguish conjunctions and disjunctions, and there are no true branches and false branches: these aspects are determined, for each node individually, in the acceptance conditions associated with the node (although the partitioning of L does indicate the effect of a node on its parents).

We will express the acceptance conditions as in [2] and as given for the Wild Animals cases in section 1.1: that is, as a series of tests, to be applied in turn (until one succeeds), and a default. There are essentially three patterns. Consider Issue1 and its children AF1 and AF2. These may form a conjunction:

AC1

```
Accept Issue1 if AF1 and AF2.
Reject Issue1.
```

Alternatively they may represent independent tests, effectively a disjunction, giving rise (as in logic programs) to two clauses and the default:

AC2

```
Accept Issue1 if AF1.
Accept Issue1 if AF2.
Reject Issue1.
```

Finally the two abstract factors may have different effects on the acceptability of their parent, giving rise to an exception or rebuttal structure:

AC3

```
Reject Issue1 if AF2,
```

```
Accept Issue1 if AF1.
Reject Issue1.
```

In AC3, AF2 is the exception. Only when AF2 does not hold, does AF1 suffice to make Issue1 acceptable. Note that the defaults may either suggest acceptance or rejection. The order of the tests represents the relative priority of the two abstract factors. Thus for example Issue2 might have acceptance conditions:

AC4

```
Accept Issue2 if AF3.
Reject Issue2 if AF4.
Accept Issue2 if AF5.
Reject Issue2.
```

which would represent a preference for AF3 over AF4, and for AF4 over AF5.

In the previous section we saw that there were essentially four types of change:

- Broadening: called *genralisation* in [45] in which a concept is modified to apply to more cases;
- Narrowing: called *restriction* in [45] in which a concept is modified to apply to fewer cases;
- Priority: called *relevance* in [45], in which an existing preference relation is changed; and
- Effect: called *inversion* in [45], in which a concept changes the party it favours.

In the following subsections we will look at each of these in turn.

3.1 Broadening

Whereas [45] envisaged changes being reflected in the structure of their decision tree, theory change in an ADF will not always result in a change in the visible structure, since sometimes the effect can be achieved by modifications to the acceptance conditions. For example, consider AC1. We can broaden Issue1 by amending the conjunction to a disjunction, which would modify it to give AC2. Such an example is provided by the automobile exception to the US Fourth Amendment. In *California v Carney*[1] all are agreed that the automobile exception arises from the greater urgency arising from the mobility of cars, and from the reduced expectations associated with vehicles, since they are subject to routine inspection. There is, however, some discussion about whether these factors can be considered independently (AC2) or whether both are needed for the exception to apply (AC1).

Sometimes, however, a change in the structure is required. There is no discussion of privacy in the original automobile exception case (*Carroll v United States*[2]), although the reduced expectations of privacy associated with automobiles is an established feature of such cases by the time of *South Dakota v Opperman*[3]. Thus at sometime between 1925 and 1976, the reduced expectations of privacy must have been introduced. Whether this broadened or narrowed the concept depends on whether AC1 or AC2 is the appropriate representation of the new acceptance conditions: for broadening, the new factor will give rise to an extra disjunctive test (AC2). The important difference is that the new node will introduce a new *value* (in this case privacy) into consideration of the issue, reflecting the key role given to purposes in [22].

[1] 471 U.S. 386 (1985)
[2] 267 U.S. 132 (1925)
[3] 428 U.S. 364 (1976)

If the new node is at the Issue level, it will require elaboration in terms of abstract and base level factors, Thus when Issue1 was introduced to the ADF of Figure 1, it would bring with it AF1, AF2, BL1, BL2 and BL3 (unless BL3 was already present because of its role in determining AF3). Similarly, if the new node is an abstract factor, it may require new base level factors to decide its acceptance. Finally the new node may be a base-level factor, which will not require anything additional in the ADF, although, if it cannot be seen as a point of an existing dimension, it will require a new dimension, if we are using dimensions. Typically a new value will give rise to a new dimension, which may further lead to the identification of additional base level factors, for use in later cases. It is, however, possible, that (as with BL3 above) no new nodes are needed since the node already exists in another part of the ADF. In such a case we need a new link, but no new node. Note that to broaden the concept the node must appear as a disjunctive test in the acceptance condition of the new abstract factor.

Broadening can also be effected by removal of a node. For example if an exception is removed, the concept will apply to more cases. For instance, removing AF4 from AC4 above will broaden Issue 2. Also broadening is effected if the node removed was used conjunctively in the acceptance conditions of the parent. For example in AC1, if AF2 is removed, this will broaden Issue1 so that it applies to all cases with AF1, rather than just the subset for which AF2 also holds.

3.2 Narrowing

Narrowing is similar to broadening but relies on the introduction of conjunctions rather than disjunctions. Thus changing from AC2 to AC1 will narrow Issue1. Similarly a new node will narrow a concept if it appears in the truth conditions either as an additional conjunct on one or more tests, or as an exception in the manner of AF2 in AC3.

Narrowing can also be effected by removing a node which was appearing in a disjunctive test (e.g. AF1 or AF2 in AC2). This, however, is quite a drastic change: first removing a node (if it is higher than the base level factors) will prune off a whole subtree, so that a number of factors which were previously considered relevant will cease to be so. This may also result in the disappearance of one or more dimensions, which will radically change the case law of the domain. Moreover, the removal of a node may result in the removal of a value, which might need some kind of societal shift to provide a justification.

There are examples of a value being rejected. For instance, consider the opinion of Brennan J. in *Furman v Georgia*[4], a US Supreme Court case involving a decision as to whether capital punishment should be considered "cruel and unusual" under the Eighth Amendment. The case is discussed in [11]. Among the questions raised was whether retribution should still be considered a legitimate purpose of punishment in the US of the seventies. Although it was not agreed that retribution had ceased to be a worthy value, Brennan holding that any retributive needs could be satisfied by a fate short of death, the Court could have held that this was so, requiring that nodes justified by this value be removed from the ADF. The problem is that such a move would represent a clear departure from *stare decisis*. In the UK, the House of Lords had been bound to follow its own previous decisions since *London Street Tramways v London County Council [1898] AC 375*. This obligation was, however, removed in a *Practice Statement* of 1966, which enabled the House of Lords to adapt English law to meet changing social conditions. This power was exercised in *R v G and R*[5], in which the

[4] 408 U.S. 238 (1972)

[5] R v G and R [2003], AC 1034

House of Lords overruled its decision in *Caldwell*[6]. These cases concerned "recklessness". This power is rarely exercised: there were fewer than twenty such instances in the first forty years of its operation (some eighty or ninety were heard in each of these years).

The US Supreme Court is able to overrule itself, but is also reluctant to do so. With respect to the automobile exception, the case of *California v Acevedo*[7] seemed to overrule both *US v Chadwick*[8] and *Arkansas v Sanders*[9] by holding that a warrantless search of an automobile was permissible, even when the probable cause only extended to a container or item of luggage within the vehicle; Blackmun J states:

> Although we have recognized firmly that the doctrine of *stare decisis* serves profoundly important purposes in our legal system, this Court has overruled a prior case on the comparatively rare occasion when it has bred confusion or been a derelict or led to anomalous results.

Thus removal of a node is possible, but not to be undertaken lightly. Indeed, Blackmun claims that *Avecedo* did not in fact overrule the previous cases. This, however, was disputed in the dissent by Stevens J with Marshall J concurring:

> Relying on arguments that conservative judges have repeatedly rejected in past cases, the Court today – despite its disclaimer to the contrary – enlarges the scope of the automobile exception

The discussion in this section indicates that removal of a node is possible, but is something of a last resort. Very often a similar effect can be achieved by reordering the tests in the acceptance conditions. Thus in AC3, the absence of AF2 is a necessary, but insufficient condition for the acceptance of Issue1. If it were placed below AF1, its presence would become a sufficient (but unnecessary) condition. In fact, following ths modification, it would have no effect, since the default would lead to rejection of Issue1 anyway if it were absent. Retaining it, however, allows for the possibility of an additional node for acceptance being introduced at a later date, which if placed below AF2 would revive its usefulness in some cases. This leads us to consider priorities.

3.3 Priority

The ordering of the tests in the acceptance conditions is determined by the relevant priorities given to the children nodes. These can be explained in terms of value preferences [1]. Note that the priorities are organised on a node-by-node basis, and there is no requirement for consistency of preferences across nodes. This constraint could be imposed, but we currently feel that the flexibility has more benefits than costs.

Thus a change in priorities is effected by reordering the tests. AC4 reflects the ordering $AF5 \prec AF4 \prec AF3$, so if AF4 became preferred to AF3, this could be effected by exchanging the position of the corresponding tests.

Changes in priority are likely to be quite common. Firstly it is very likely that not all the priorities represented in the ADF will have a justification in the precedent cases, but will have been determined by the judgement of the knowledge engineer. Whilst, as more cases are decided, many of the judgements will be confirmed, some new cases may require

[6] R v Caldwell [1982] AC 341
[7] 500 U.S. 565 (1991)
[8] 433 U.S. 1 (1977)
[9] 442 U.S. 753 (1979)

tuning of the acceptance conditions for particular nodes. Moreover, and more importantly, it needs to be recognised that the value preferences of society change over time. Sometimes this will be a case of evolution, as with *retribution* discussed above for capital punishment, but sometimes it will be more like a pendulum. This is well described in [29], particularly in the context of values expressed as legal principles. These often come in opposing pairs: for example it is desirable to draw a bright line in the interests of legal certainty, but it also necessary to deal out justice in the light of the nuances of the facts of individual cases. Similarly the amount of discretion it is held proper for judges to exercise will vary. As the pendulum swings, so the preference expressed in the associated case law may change. For example, although it was held in *Pierson v Post*, as discussed in [21] that the need for clarity outweighed any social benefit that would arise from encouraging fox hunting, the balance may have been struck differently at another time or in another place.

3.4 Effect

The last category of change in [45] concerns the effect of a factor. In CATO factors are supposed to always favour one party or the other, and this is an inherent feature of the factor. This is a key part of factor based reasoning, and an essential component of both [18] and [37]. Thus in the discussion of *Pierson v Post* in [21] we have a factor *NotCaught* which definitely favours the defendant. Much of the argument, however, concerns what should count as caught and, as discussed in [17], the whole case may turn on this point. In *Pierson*, it is argued, following the authority of Justinian, that actual bodily possession was required. It was, however, also argued, following other authorities, that mortal wounding, or even certain capture, would suffice. The counsel for Post argued that even being in hot pursuit should count (see the discussion in [17]). Here we might see *NotCaught* as an abstract factor and *BodilyPossession*, *MortalWounding*, *CertainCapture* and *HotPursuit* as base level factors. The presence of any of the last three of these would lead to the acceptance of *NotCaught* and so all three would be regarded as pro-defendant factors. Only *BodilyPossession* is pro-plaintiff.

AC5

```
Accept NotCaught if HotPursuit
Accept NotCaught if MortalWounding
Accept NotCaught if CertainCapture
Reject NotCaught if BodilyPossession
Accept
```

Now if later it was decided that Justinian was too restrictive, then one of the other factors would change from accept to reject. This is more than a re-ordering of priorities: it involves a change in the party favoured by the factor, to reflect some social change. As the pendulum discussed above in connection with priorities swings, it will move along the dimension and modify the cross-over point, turning pro-plaintiff base level factors into pro-defendant factors as it moves one way, and back again when it returns. Essentially we are modifying what we are prepared to accept as promoting the value underlying the dimension.

3.5 Importance of Values

What the above shows is how right Berman and Hafner were to link changes in case law to dimensions and changes in recognised social values and their priorities in [22]. Although some modifications, especially in the early stages of case law development, will draw on values already present, many of the modifications, of all four types, will result from values becoming accepted or rejected, the preferences between them being modified, or changes in what counts as promotion. This is how the law can reflect social change, while retaining the same legislation.

4 Berman and Hafner's Red Flags

In the previous section we have considered how changes over time can be incorporated in the representation of a body of case law as an ADF. Our perspective thus far has been very much *post hoc*. In [22], however, Berman and Hafner were especially concerned with *anticipating* the need for change. In particular they wanted to know how confident they should be in a prediction based on a representation of existing case law. To this end they identified five "red flags" which were supposed to warn the user of a putative system to be cautious about relying on the predictions, since the law might be about to change, and the current case might be the one where that change was to become manifest.

The first of these red flags was where the precedent had been explicitly overruled. It is quite possible that the decision not to modify the representation immediately had been taken by the user, since a single decision might be a rogue decision, and we might find it ignored in subsequent decisions. The existence of such a decision should, however, require that the system be used with caution until it became clear whether the change was genuine or a mistake. What should be done, during the period when change is under consideration, is to annotate the program so that the user is made aware that the prediction offered relies on a decision which has been overruled.

An explicit overruling is easy enough to spot, provided someone is monitoring the decisions, but the second red flag relates to implicit overruling, and so needs more than a superficial look at the decision. If, however, we have a working system to maintain, the system itself can be used to detect such decisions. If we run the case through the system and a divergent answer is given, we may suspect an implicit overruling. The node which gave rise to the divergence can be readily spotted as described in [2]. We can then annotate the program, as for an explicit overturning, to sound a note of caution to users until it becomes clear whether we should adopt the change.

The third red flag is where the *ratio decidendi* is overruled, although the facts differ significantly. In [22] the example was where an issue is applicable to two distinct domains, but was resolved differently in the two domains[10]. Clearly this difference in value preferences should ring a warning bell, but it will be hard to detect: we may not be monitoring cases in this other domain, let alone have an operative system which will detect the divergence. A problem with representing law is that to build a software system we need to scope the domain: indeed scoping the system and deciding what will be in the system and what will not is an essential early step in Software Engineering methodologies (e.g. [31]). But the law

[10] The issue in [22] is whether a case should be tried in the jurisdiction where the event took place, or in the jurisdiction which was the "centre of gravity". Thus if two New Yorkers have a motor accident in New Hampshire, where should the case be heard? Similar cases could arise in other domains: for example if two New Yorkers struck an agreement when they happened to be in New Jersey.

is, in the phrase associated with Ronald Dworkin, a "seamless web", and such scoping may be difficult or even impossible. If detected, this red flag can be handled in the same way of the previous two, but detection cannot be assured.

The fourth red flag is when a "shift in the relative priority of competing purposes is in evidence from the court's tendency to make a rule increasingly narrow by "distinguishing away situations where one might have expected the rule to apply", [22] (page 45). We saw in section 3.2 how a concept could be narrowed, by adding conjuncts to truth conditions, and by adding exception structures. In order to pick up on this flag a person must notice a pattern in the series of changes, namely that the role of nodes related to a particular value is being restricted by nodes relating to some other value, or values. Such a pattern is likely to indicate a loss of favour of the value being restricted. An alternative pattern might arise from nodes relating to a particular value being used to restrict nodes relating to other values, suggesting that the preference for the value is increasing. Such patterns will, however, emerge only from a reasonably large number of modifications arising in the domain of interest. Probably the clearest sign is when we get a new exception structure, of the form of AC3. AC3 shows a clear preference for the value associated with AF2 over that associated with AF1, and may give us warning that a shift in value preferences has occurred. This in turn should lead us to reconsider other expressions of preference and where the new preference is not respected annotate it with an appropriate warning.

The final red flag is where "a general shift in the relative priority of competing purposes is in evidence from cases across legal domains" [22] (page 46). Like the third red flag this requires the monitoring of domains other than the one represented. And like the fourth red flag it is somewhat nebulous. The chances of detection are thus lower than the other red flags. Perhaps, however, to talk about competing purposes in evidence from legal decisions is to place too narrow bounds on where such shifts in values can be discovered. Value preferences are manifested in social attitudes as well as law and while sometimes the law is ahead of the public in terms of social change (capital punishment is a UK example), in others it lags behind. An example of the latter can be perhaps been seen in the area of rights over a matrimonial home as seen in *Pettitt v Pettitt*[11]. We consider that identifying instances of this fifth red flag might prove difficult and unreliable

In [22] an algorithm is provided for using these red flags to provide a numeric measure to which a precedent had been weakened by subsequent red flags. This, however, does not discuss how red flags can be identified and so has not really been used in any empirical investigation, and so we will not discuss it further.

5 Related Work

In this section we will discuss some related work.

5.1 Detecting Change

The investigation in [46] concerned the response to surprising decisions which apparently went against existing case law: so-called *Black Swans*: "novel, surprising, provocative, ex-

[11] Pettitt v Pettitt [1970] AC 777. Pettitt is an English case in which the wife had used her own money to buy a house during the marriage, meaning the title to the house was in the wife's name, and both she and her husband lived there until the wife left the husband. The husband claimed that he had carried out a considerable number of improvements to the house and garden and sought a beneficial interest in the proceeds of sale of the property.

ceptional cases". In order for Black Swans to represent a change rather than a single aberration they need to be followed by Gray Cygnets, cases which reinforce the decision and accelerate legal change. The domain investigated was recovery of damages by a remote buyer, a sequence of cases initiated by the decision in *Thomas and Wife v. Winchester*[12]. Like Berman and Hafner's red flags, this work relates to when we need to change our ADF, rather than any particular changes we might need to make.

5.2 Responding to Change

Although there has been work on detecting, anticipating and responding to change, little attention has been paid to the evolution of case law. Exceptions are [36] and [43], but neither of these seem to have attracted any follow up attention. Perhaps this is a topic worthy of future AI and Law investigation. The maintenance of knowledge based systems, always something of a Cinderella subject, seems to have attracted little attention since the 1990s, and so [30] remains a useful reference.

5.3 Theory Construction

In both [2] and [18] the intention is to encapsulate a theory derived from a body of case law. But whereas [2] considers all the precedents to be included at the outset, and aims to capture them all in the ADF, in [18] the theories include cases one by one and build upon them by adding more cases from the available precedents to refine the theory until the domain is explained as far as it can be. Therefore [18] also comes with a set of operators for extending the theory, as new cases are included. Although these are not new cases, but cases chosen from the case background, there are clear similarities with the changes to a theory expressed as an ADF as described in section 3.

A theory in [18] is a five-tuple comprising a set of cases, a set of factors, a set of rules, a set of preferences between these rules, and a set of value preferences intended to justify the rules. As an example we will consider the theory intended to explain *Pierson v Post*[13], *Keeble v Hickeringill*[14] and *Young v Hitchens*[15], the three wild animals cases at the heart of [21]. The eventual theory that emerges in [18] is:

Cases: Pierson, Keeble, Young.
Factors: NotCaught, PlaintiffEarningLivelihood, DefendantEarningLivelihood, PlaintiffOwnedLand.
Rules:
Rule1: NotCaught $\rightarrow$ Defendant.
Rule2: PlaintiffEarningLivelihood $\rightarrow$ Plaintiff.
Rule3: DefendantEarningLivelihood $\rightarrow$ Defendant.
Rule4: NotCaught $\wedge$ DefendantEarningLivelihood $\rightarrow$ Defendant.
Rule Preferences: Rule2 $\succ$ Rule1. Rule4 $\succ$ Rule2.[16]
ValuePreferences: Productivity $\succ$ Clarity. [Productivity,Clarity] $\succ$ Productivity.

[12] 6 N.Y. 397 (1852).
[13] 3 Cai. R. 175, 2 Am. Dec. 264 (N.Y. 1805)
[14] 103 ER 1127 (1707)
[15] 6 QB 606, 115 ER 228 (1844)
[16] There is no precedent to tell us the priority of Rule3.

Now let us view this an an ADF. First note that the structure is much simpler than our example in Figure 1: all the factors are base level factors, and there is just one root node above them (*Decision*). The rules of [18] are all just one step and no chaining is used. Thus all the rules must be placed in the acceptance condition of *Decision*. These acceptance conditions will be:

AC6

```
Defendant if DefendantEarningLivelihood and NotCaught.
Plaintiff if PlaintiffEarningLivelihood.
Defendant if NotCaught.
Defendant.
```

Value preferences are not explicit in the ADF, but we could associate them with rules: Rule1 promotes Clarity, Rules 2 and 3 promote Productivity and Rule 4 promotes both. Thus the ordering in AC6 follows the value preferences of the theory. The result is a much simpler structure, but one which lacks many of the software advantages of a more elaborate ADF as given in section 1.1. In particular the rule base is monolithic, lacking the modularisation typically provided by an ADF. Moreover the flatness of the structure inhibits the transparency of the theory, and the explanation capabilities of an implemented system. The opacity of explanations produced from the theories of [18] can be seen from [28].

Turning now to the theory construction operators of [18], the first is *includeCase*. This is not needed in the ADF, which has the complete set of cases from the outset. The second, *includeFactor*, corresponds to adding a node to the ADF. Whether it broadens or narrows the concept to which it is added depends on the party favoured by the factor, and the role played by the new factor in the acceptance condition. Note too that the factor brings with it a value, and we saw that adding a node was often motivated by the desire to introduce a new value into consideration. The third is *factorMerging* which brings together two factors to act as a conjunctive antecedent. This is similar to narrowing a concept in an ADF by, e.g. moving from AC2 to AC1. The fourth constructor is *ruleBroadening*, which corresponds to broadening in the ADF, as when we move from AC1 to AC2. The fifth operator is *PreferenceFromCase*, which corresponds to a change in rule order in an acceptance condition, as discussed in relation to priorities in section 3.3. above. The sixth operator is *RulePreferenceFromValuePreference*. This also is effected in the ADF by a change in the order of the test in an acceptance condition, but the motivation differs. The same is true of the remaining operators, *ArbitraryRulePreference* and *ArbitraryValuePreference*, except these are not motivated by a case. Thus there are clear correspondences between the operators of [18] and the ways in which an ADF is modified. One difference is that [18] is much more explicit about the sources of change, and the values. This suggests that we should consider annotating the tests of an ADF with the cases they derive from and the values they promote. In this way we would combine some of the strengths of the theory construction approach with the software engineering advantages of using an ADF. Moreover such annotations would represent the sort of documentation which should accompany any software design.

6 Concluding Remarks

In this paper we have drawn inspiration from [22] which considers the temporal context of a system based on legal cases. That paper itself develops ideas relating to purpose put forward in [21]. This is only natural: while technological developments can make changes in the law necessary, as when the invention of cinema necessitated a rethinking of copyright law, and as

perhaps robots and autonomous agents will necessitate changes in the near future, the main driver for legal change remains the social context in which the law operates. The law has a purpose, but this is closely related to the currently favoured social values. Sometimes social change can be seen as undisputed progress, as with the rights of women in the last century, and attitudes to same sex marriage in this century, from which (it is to be hoped) there is no going back. In other changes in social attitude, such as the degree of judicial discretion acceptable, the pendulum swings first one way then another, in accordance with intellectual fashion.

Whatever the source of change, it must be accommodated in the representation of case law domains, and here we have discussed how change can be accommodated in a current proposal for an engineering methodology for such a system, and related this account to previous work on change.

Perhaps the most interesting feature of [22] is its stress on dimensions. Although dimensions were a feature of legal case based reasoning in HYPO, more than thirty years ago, they have remained largely dormant while reasoning with factors has been fully explored. But in recent years dimensions have been revived as an essential bridge between factors and the facts describing a case from the perspective of the real world, as in [42] and [13]. It is this kind of prescience that makes the work of Carole and Don on case based reasoning in law still relevant after more than twenty years.

Acknowledgments

Since this paper is to appear in a special issue celebrating the life and work of Carole Hafner, we feel that we should acknowledge Carole's contribution to AI and Law research in general and to our research in particular. The debt of the second and third authors to the work of Carole and Don Berman, in particular [21], is considerable: much of our work has explored value-based reasoning in case law, explicitly building upon the teleological aspects of reasoning with cases identified there. Carole and Don were an all too rare example of an equal partnership of AI and Law, with Carole supplying the AI and Don the law. The balance between the disciplines is meticulous: there is a real legal problem and a well thought through practical, computable, solution to that problem. This contrasts with much AI and Law work: all too often we find either a real legal problem and some hand-waving suggestions of how it might be solved, or a computer solution in search of a problem, and the example is at best a toy problem and at worst purely symbolic. Carole and Don showed us what can be done when the two disciplines work together in harmony. If she had done nothing else - and she did a lot of other valuable work, especially relating to conceptual retrieval and ontologies, not to mention her important role in creating an AI and Law community - Carole's contribution to AI and Law in the early nineties papers with Don would represent one of the more significant achievements in the field. After more than twenty years of study, these papers continue to yield fresh insights.

References

1. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Factors, issues and values: Revisiting reasoning with cases. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 3–12. ACM, 2015.

2. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artificial Intelligence and Law*, 24(1):1–49, 2016.
3. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Angelic secrets: Bridging from factors to facts in US trade secrets. In *Proceedings of Jurix 2016*, In Press, 2016.
4. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Statement types in legal reasoning. In *Proceedings of Jurix 2016*, In Press, 2016.
5. V. Aleven. *Teaching case-based argumentation through a model and examples*. PhD thesis, University of Pittsburgh, 1997.
6. M. Araszkiewicz, A. Łopatkiewicz, A. Zienkiewicz, and T. Zurek. Representation of an actual divorce dispute in the parenting plan support system. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 166–170. ACM, 2015.
7. K. D. Ashley. *Modeling legal arguments: Reasoning with cases and hypotheticals*. MIT press, 1991.
8. K. D. Ashley and S. Brüninghaus. A predictive role for intermediate legal concepts. In *Proceedings of JURIX*, pages 153–162, 2003.
9. K. D. Ashley, C. Lynch, N. Pinkwart, and V. Aleven. A process model of legal argument with hypotheticals. In *JURIX*, volume 189, pages 1–10, 2008.
10. K. Atkinson and T. Bench-Capon. Legal case-based reasoning as practical reasoning. *Artificial Intelligence and Law*, 13(1):93–131, 2005.
11. T. Bench-Capon. Towards computational modelling of supreme court opinions: Furman v Georgia. In *K. Atkinson (ed) Modelling legal cases*, pages 329–42. IDT Series, Universitat Autònoma de Barcelona Barcelona, 2009.
12. T. Bench-Capon. Representing Popov v Hayashi with dimensions and factors. *Artificial Intelligence and Law*, 20(1):15–35, 2012.
13. T. Bench-Capon and F. Bex. Cases and stories, dimensions and scripts. In *Legal Knowledge and Information Systems: JURIX 2015: The Twenty-Eighth Annual Conference*, volume 279, page 11. IOS Press, 2015.
14. T. Bench-Capon and F. Coenen. Exploiting isomorphism: development of a KBS to support british coal insurance claims. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 62–68. ACM, 1991.
15. T. Bench-Capon and T. F. Gordon. Isomorphism and argumentation. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, pages 11–20. ACM, 2009.
16. T. Bench-Capon, H. Prakken, A. Wyner, and K. Atkinson. Argument schemes for reasoning with legal cases using values. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pages 13–22. ACM, 2013.
17. T. Bench-Capon and E. L. Rissland. Back to the future: dimensions revisited. In *Proceedings of JURIX 2001*, pages 41–52, 2001.
18. T. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1):97–143, 2003.
19. T. J. M. Bench-Capon. The missing link revisited: The role of teleology in representing legal argument. *Artificial Intelligence and Law*, 10(1-3):79–94, 2002.
20. D. H. Berman and C. D. Hafner. Incorporating procedural context into a model of case-based legal reasoning. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 12–20. ACM, 1991.
21. D. H. Berman and C. D. Hafner. Representing teleological structure in case-based legal reasoning: the missing link. In *Proceedings of the 4th international conference on*

Artificial intelligence and law, pages 50–59. ACM, 1993.
22. D. H. Berman and C. D. Hafner. Understanding precedents in a temporal context of evolving legal doctrine. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 42–51. ACM, 1995.
23. L. K. Branting. Reasoning with portions of precedents. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 145–154. ACM, 1991.
24. P. Bratley, J. Frémont, E. Mackaay, and D. Poulin. Coping with change. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 69–76. ACM, 1991.
25. G. Brewka, S. Ellmauthaler, H. Strass, J. P. Wallner, and S. Woltran. Abstract dialectical frameworks revisited. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 803–809. AAAI Press, 2013.
26. G. Brewka and S. Woltran. Abstract dialectical frameworks. In *In Proc. KR-2010*, pages 101–111, 2010.
27. S. Bruninghaus and K. D. Ashley. Predicting outcomes of case based legal arguments. In *Proceedings of the 9th international conference on Artificial intelligence and law*, pages 233–242. ACM, 2003.
28. A. Chorley and T. Bench-Capon. Agatha: Using heuristic search to automate the construction of case law theories. *Artificial Intelligence and Law*, 13(1):9–51, 2005.
29. G. Christie. *The notion of an ideal audience in legal argument*. Springer Science & Business Media, 2012.
30. F. Coenen and T. Bench-Capon. *Maintenance of Knowledge-Based Systems: Theory, Techniques & Tools*. Academic Press, 1993.
31. T. M. Connolly and C. E. Begg. *Database Solutions: A step-by-step guide to building databases*. Pearson Education, 2004.
32. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357, 1995.
33. C. D. Hafner. Conceptual organization of case law knowledge bases. In *Proceedings of the 1st international conference on Artificial intelligence and law*, pages 35–42. ACM, 1987.
34. C. D. Hafner and D. H. Berman. The role of context in case-based legal reasoning: teleological, temporal, and procedural. *Artificial Intelligence and Law*, 10(1-3):19–64, 2002.
35. C. D. Hafner and M. Lauritsen. Extending the power of automated legal drafting technology. In *JURIX*, volume 165, pages 59–68, 2007.
36. J. Henderson and T. Bench-Capon. Dynamic arguments in a case law domain. In *Proceedings of the 8th international conference on Artificial intelligence and law*, pages 60–69. ACM, 2001.
37. J. F. Horty and T. J. M. Bench-Capon. A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20(2):181–214, 2012.
38. E. H. Levi. An introduction to legal reasoning. 1961.
39. L. Lindahl and J. Odelstad. Intermediate concepts in normative systems. In *Deontic Logic and Artificial Normative Systems*, pages 187–200. Springer, 2006.
40. N. F. Noy and C. D. Hafner. The state of the art in ontology design: A survey and comparative review. *AI magazine*, 18(3):53–74, 1997.
41. H. Prakken. An exercise in formalising teleological case-based reasoning. *Artificial Intelligence and Law*, 10(1-3):113–133, 2002.

42. H. Prakken, A. Wyner, T. Bench-Capon, and K. Atkinson. A formalization of argumentation schemes for legal case-based reasoning in aspic+. *Journal of Logic and Computation*, 25(5):1141–1166, 2015.
43. A. Priddle-Higson. *Computational models of ontology evolution in legal reasoning*. M. Phil Thesis: The University of Edinburgh, 2010.
44. E. L. Rissland. Dimension-based analysis of hypotheticals from supreme court oral argument. In *Proceedings of the 2nd international conference on Artificial intelligence and law*, pages 111–120. ACM, 1989.
45. E. L. Rissland and M. T. Friedman. Detecting change in legal concepts. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 127–136. ACM, 1995.
46. E. L. Rissland and X. Xu. Catching gray cygnets: an initial exploration. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 151–160. ACM, 2011.
47. G. Sartor. Teleological arguments and theory-based dialectics. *Artificial Intelligence and Law*, 10(1-3):95–112, 2002.
48. M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. The British Nationality Act as a logic program. *Communications of the ACM*, 29(5):370–386, 1986.
49. D. B. Skalak and E. L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial intelligence and Law*, 1(1):3–44, 1992.
50. A. Wyner and T. Bench-Capon. Modelling judicial context in argumentation frameworks. *Journal of Logic and Computation*, 19(6):941–968, 2009.
51. A. Wyner, T. Bench-Capon, and K. Atkinson. Arguments, values and baseballs: Representation of Popov v. Hayashi. In *Proceedings of JURIX 2007*, pages 151–160, 2007.
52. A. Wyner and R. Hoekstra. A legal case owl ontology with an instantiation of Popov v. Hayashi. *Artificial Intelligence and Law*, 20(1):83–107, 2012.