Factors, Issues and Values: Revisiting Reasoning with Cases

Latifa Al-Abdulkarim
Dept. of Computer Science
University of Liverpool
UK
latifak@liverpool.ac.uk

Katie Atkinson
Dept. of Computer Science
University of Liverpool
UK

Trevor Bench-Capon
Dept. of Computer Science
University of Liverpool
UK

ABSTRACT

In this paper we revisit reasoning with legal cases, with a view to articulating the relationships between issues, factors, facts and values, and to identifying areas for future work on these topics. We start from the different ways in which attempts have been made to go beyond *a fortori* reasoning from the precedent base, so that conclusions not fully justified by the precedents can be drawn. We then use a particular example domain taken from the literature to illustrate our preferred approach and to relate factors and values. From this we observe that much current work depends critically on the ascription of factors to cases in a Boolean manner, while in practice there are compelling reasons to see the presence of factors as a matter of degree. On the basis of our observations we make suggestions for the directions of future work on this topic.

1. INTRODUCTION

Modelling reasoning and arguing with legal cases has been a central concern of AI and Law since its very beginnings [34], [27]. Over time, a view has developed that legal reasoning has to progress in stages. First, facts must be determined on the basis of evidence (e.g. [18]). Next these facts must be used to ascribe legally significant predicates to the case. These predicates, which serve as intermediaries between the world of fact and the world of law (e.g. [21]), have been termed intermediate concepts [33], [8], but have been more often called *factors*, following the highly influential HYPO [6] and CATO [4] projects[1], which are still foundational for discussion of reasoning with legal cases in AI and Law. Factors provide a level of abstraction that enables cases with different fact situations to be compared, matched and distinguished, and also have implications for the outcome of the cases. The set of factors associated with a domain are either present in, or absent from, a case in that domain, and the presence of a particular factor supports a decision for one of the parties to the case, either the plaintiff or the defendant. Given a set of precedent cases described in terms of the factors present in them, the idea is that in order to determine the arguments applicable to a new case, the new case should be matched with the precedents and good matches will provide applicable arguments. Neither HYPO nor CATO predict decisions, but the legal principle of *stare decisis* enables the precedents to offer reasons for deciding the new case in a particular way: this feature was used to make predictions in IBP [23].

In HYPO and CATO the precedents are deployed through a process of argumentation (usually called *three ply argument*), in which a matching precedent is cited in the first ply, attempts are made to distinguish the case or provide counter examples in the second ply, and in the third ply the objections are themselves subject to a rebuttal. Any decision is left to the users who may choose to ascribe legal consequences, depending on whether they believe that, in the light of the arguments, the particular set of factors in the new case favour the plaintiff or the defendant. Most attention in AI and Law has been paid to hard cases, which have been appealed to a higher court. In these appeals the facts are usually taken as already established by the lower court: for this reason the majority of AI and Law systems subsequent to HYPO and CATO, such as [37], [15] and [10] have been able to consider cases as sets of factors, and have concentrated on how an outcome can be reached on the basis of arguing about the competing merits of the plaintiff and defendant factors.

A major challenge for the above account is how to decide cases which do not have a precedent that controls the case. There can be a reasonably large number of factors (26 in the case of CATO[2], 13 favouring the plaintiff and 13 favouring the defendant). This means that there are 2^{13} (8192) factor combinations for each side, and hence potentially 2^{26} (67,108,864) comparisons of sets of plaintiff factors with sets of defendant factors. In contrast, the number of cases used in these systems is small: HYPO used 33 cases [6], CATO had a case base of 148, extended to 186 for the Issue Based Prediction system (IBP) [23], and AGATHA [24] (another system based on the same domain that will be discussed in this paper) used 32. This set of 32 cases was also used recently to evaluate an approach [2] based on Abstract Dialectical Frameworks [22]; the evaluation can be found in [3]. Given this discrepancy between the number of potential comparisons and the number of precedents[3]

[1]HYPO predominately uses *dimensions* rather than factors, although factors were used in the later stages of the project and are discussed in [6]. Factors were the basis of CATO, which did not use dimensions, and it is CATO's notion of factors that we will use here. For a discussion of the relation between dimensions and factors see [39].

ICAIL '15, June 08 - 12, 2015, San Diego, CA, USA

http://dx.doi.org/10.1145/2746090.2746103.

[2]We will use CATO as the the predominant source of our illustrations, as it is well known and documented and offers a clear account of reasoning with factors, which has been used as the basis of discussion in later work such as [37], [15], [24], [30] and [3]. CATO operates in the domain of US Trade Secret law.

[3]The problem is exponential in the number of factors. Exponential problems are usually associated with NP-completeness (a form of computational intractability) in computer science, but here the

available to justify them, it is essential to have an approach to deciding comparisons, which takes us beyond what is fully justified by the case base of precedents.

In CATO itself, this is achieved by organising the 26 base level factors into a factor hierarchy. Each base level factor is seen as the child of one or more *abstract factors*, and acts as a reason for the presence or absence of its parent(s). Abstract factors also favour the plaintiff or the defendant. If a factor favouring a party to the case is present in the precedent but absent in the new case, a factor with the same parent present in the new case may be substituted for it when arguing for that party. (Similarly, *mutatis mutandis*, when arguing for the other side.) This is known as *downplaying* the distinction. If such substitution is not possible, the distinction may be *emphasised*. Such substitution allows us to decide comparisons not available in the precedents. In CATO there are 11 abstract factors (7 favouring the plaintiff and 4 favouring the defendant), and 5 *issues*, which provide a layer of abstraction above the abstract factors. Of these issues, 4 favour the plaintiff and 1 the defendant. If we consider cases as sets of abstract factors we reduce the potential comparisons to 2^{11} (2048). Although this is a significant reduction, it remains far in excess of the size of the precedent base in existing systems. Considered in terms of issues alone, however, the comparisons reduce to 2^5 (32), which, perhaps co-incidentally, was the number of cases used in HYPO and AGATHA.

An alternative approach to limiting the comparisons was suggested in [16], most fully developed in [15] and empirically evaluated in [24] and [25]. The idea here is that the factors can be related to social purposes (usually termed *values* in the literature), and that comparison of factors is determined by preferences between the purposes or values they promote. The preferences may change over time and may differ between jurisdictions. On this account the precedents determine value preferences and these in turn can be used to decide comparisons not present in the precedent base. This reduces the number of comparisons further (five values were used in [24]: the same values may apply to factors favouring both parties and are themselves strictly ordered). Values have since been used in argumentation scheme approaches (e.g. [13] and [10]), in which the value preferences can prioritise the rules and decide the winner when arguments conflict. In these accounts values attach to a rule through their association with factors, but another recent discussion of values in case based reasoning [5] argues that not only may rules promote purposes, but also the particular conditions in the rules may each have their own purposes which justify their inclusion in the rule. Other recent discussions of the role of values in reasoning with legal cases include [29].

Although much work has been done using values, another approach does not use values at all. IBP [23], [8] and [7] takes the hierarchical approach of CATO a step further by organising the issues into a logical model[4] based on the Restatement of Torts. Attempts have been made to represent the precedents as sets of rules without recourse to purposes and values, including [37] and more recently [30]. Finally the ADF approach of [2] organises issues, abstract factors and base level factors into a single structure, but does not include values in this structure.

In this paper we will discuss some of the open questions relating to factors, issues and values in legal case based reasoning. In particular we will discuss how it is possible to go beyond the comparisons established in the precedent base, and how the required comparisons should be made. In so doing we hope to establish an agenda for further research on this topic.

problem is not the number of computations, but the number of precedents we potentially need to collect to fully justify every decision. 2^{26} is likely to exceed the number of precedents in existence.

[4]A top level of logical rules with the leaves of the logical model resolved using factor based reasoning with cases was also a feature of CABARET [41].

2. WHY GO BEYOND PRECEDENTS?

When a case is brought before a court a decision has to be made: abstention is not possible. Thus where no precedent exists, the law is effectively extended, although the arguments which justify the decision are supposed to keep the case law, taken as a whole, as coherent as possible. The practical importance of going beyond the precedents is clearly shown by the empirical results of [23]. In that paper a number of systems were evaluated, including several machine learning systems, programs representing HYPO and CATO and the logical model of IBP without the case based reasoning component, as well as IBP itself. The results taken from [23] are shown in Table 1[5]. We have also included two lines representing AGATHA [24]. A number of different heuristics were investigated in [24], so we have included both a high (for brute force with the largest number of seed cases, labelled AGATHA 8) and a low figure (for the version which used A* for heuristic search). Another two lines are taken from a recent evaluation [3] of a representation using Abstract Dialectical Frameworks. One represents an ADF encoding of the CATO hierarchies[6], the other an ADF refined in the light of re-examination of the decisions of the cases wrongly decided by the direct encoding. In order to show the effect of abstention we have included an additional column showing the percentage of correct answers for which an opinion was given (akin to precision in an Information Retrieval system). Finally two hypothetical programs are included: CATO-coin and HYPO-coin, in which the abstentions are decided arbitrarily and it is assumed that half these arbitrary choices will be right and half wrong.

Table 1: Results from [23], [24] and [3]

	correct	error	abstain	accuracy	no-abst
ADF Refined	31	1	0	96.8	96.8
AGATHA 8	30	2	0	93.7	93.7
IBP	170	15	1	91.4	91.9
AGATHA A*	29	4	0	90.6	90.6
CATO-coin	163	30	0	89.0	89.0
Naive Bayes	161	25	0	86.5	86.5
HYPO-coin	152	34	0	81.7	81.7
ADF Cato	25	7	0	78.1	78.1
CATO	152	19	22	77.8	88.8
HYPO	127	9	50	68.3	93.4
IBP-model	99	15	38	72.6	86.8

What Table 1 shows is that simply guessing the outcome can improve the percentage of correct answers of most systems, although it greatly reduces the precision of, for example, HYPO, which abstains on a very large number of cases, but is correct in an excellent percentage of those for which it does yield an answer. Anything that can be done to reduce abstentions, even if the success rate is only 50%[7], is thus going to be beneficial in terms of suc-

[5]The differences in the total number of cases for CATO and IBP-model are as in [23].

[6]It should be remembered that the analysis of CATO was done to assist law students to learn how to distinguish cases, not to predict outcomes.

[7]Note, however, that the unrefined ADF seems to get the abstained cases uniformally wrong, so CATO is wise to abstain.

cessfully predicted outcomes. In the next section we will consider approaches to going beyond the precedents.

3. BEYOND PRECEDENTS

In this section we will consider three ways of going beyond precedent: rule broadening, factor reduction and partitioning.

3.1 Rule Broadening

An important landmark in modelling case based reasoning is [37] which explained how precedents could be modelled as sets of rules, thus making the logic underlying *stare decisis* more explicit. The idea was to represent precedent cases as three rules:

R1: Plaintiff factors present $\rightarrow$ Plaintiff

R2: Defendant factors present $\rightarrow$ Defendant

R3: A priority rule expressing R1 $\succ$ R2 if the plaintiff won the case and R2 $\succ$ R1 if the defendant won the case.

This captures the information in the set of precedents and can be applied to new cases, but will only be able to resolve cases already covered by precedents. Note also that there is an assumption that strengthening the antecedent is always possible, so that if a case contains a superset of the plaintiff factors, R1 will still apply. Since factors always favour the same side if present, this assumption is reasonable. There are, however, situations which suggest that we cannot do this in general: [36] has a counter example. Henry does not like to jog if it is wet or hot, but does like to jog in the rain when it is hot. Thus individually hot and wet are pro-non-jogging factors, but their combination is pro-jogging. Since we do need the ability to strengthen the antecedent, this means that we need to be careful to avoid such possibilities when identifying factors. Given the ability to strengthen the antecedent, we can say that a precedent C with pro-plaintiff factors P and pro-defendant factors D will also cover cases with a superset of P, or a subset of D or both. Thus the precedent will govern a number of comparisons as *a fortiori*, but cannot take us beyond the precedents, legitimising, for example the broadening of R1 to a case with a subset of P.

This is possible in the approach of [30]. If it is known that $R1 \succ R2$, it is possible (assuming the plaintiff wins the case) that $R4 \succ R2$, where $R4$ is $P2 \rightarrow Plaintiff$, with $P2 \subset P$. Thus it is argued in [30] that the rule for the winning side should not be the *complete* set of factors for that side, but a subset. This does make it possible to cover more comparisons, the rule broadening taking place when recording the decision rather than in response to a new case, but it does raise problems of interpretation: the appropriate subset may not be obvious. Moreover, it provides no explanation of why this particular subset of factors is used. One possibility is to use the *ratio decidendi* of the case, but this is often hard to identify and map to factors. See, however, [20] for an approach to this problem.

3.2 Factor Reduction

The problem arises from the large difference between the number of potential comparisons and the number of precedents available. Thus, if we had fewer factors, a given number of precedents would cover a greater proportion of new cases, and the need to abstain would be reduced. In the limiting case with two factors, one for each side, a single precedent would be sufficient[8]. Unfortunately we cannot simply limit the number of factors used: we need a sufficient range of factors to adequately represent the domain.

[8]In this paper we ignore the possibility of conflicting precedents. Ideally the precedent base should be free from such conflicts: when precedents are explicitly departed from or overturned they should be removed from the precedent base but we recognise that complete consistency remains unlikely.

This is motivation for the use of abstract factors in CATO. As we saw above, using the 11 abstract factors of CATO reduces the potential comparisons to 2^{11} (2048). Thus downplaying enables us to resolve cases not fully covered by precedents. Moreover, downplaying makes excellent sense in many cases. Consider cases where the secret was obtained through an illegal act, which might be bribery, theft or deception. In practice it seems to make little difference which particular illegal method was used: a trade secret obtained in any of these ways is considered misappropriated. In other cases it is perhaps less obvious that substitution should be acceptable: for example it is unclear whether disclosing the secret in a public forum should have more (or less) impact than the fact of the information being actually known to specific competitors, although both have the same abstract factor (*information known*) as parent. But perhaps the real problem is that some distinctions cannot be downplayed: for example, it is not possible (in the CATO hierarchy as given in [4]) to substitute *information known to competitors*, which relates to whether the information was a trade secret, for *information independently generated* which affects claims not to have misappropriated the information.

The approach using purposes or values is more comprehensive. Since every factor is associated with a value (in [15] at least), any case can be reduced to a comparison of values. Given five values (as in [24]), we have 32 value sets for each party, a maximum of 1024 possible comparisons. Moreover, an ordering of these values can be established using relatively few cases (see [24]). There are, however, two problems: one is that while *stare decisis* does seem to justify consistency of decisions considered in terms of factors, there is no such link between the decisions and the values. Values and purposes are often not explicitly mentioned, although such considerations are more likely to feature as one ascends through the appeals process. Second, and more seriously, there is a strong suggestion from empirical work such as [24] that different factors promote values to different degrees, and that the strong promotion of a weak value may be preferred to the weak promotion of a stronger value (see also [1]). But if we cannot use a strict ordering of values, the number of potential comparisons grows very rapidly.

One way to view these difficulties is to see them in terms of deciding how to compare two sets of values. The literature contains several suggestions, but no consensus has been reached. In [12] the following questions are posed:

1. Is promotion and demotion of values to be regarded as Boolean (either promoted or not) as in [15], ordinal (instances of promotion can be ordered), qualitative (instances of promotion can be placed into bands), as in [29], or quantitative (instances of promotion can be assigned specific numeric values) as in [24]?

2. Should we see promotion and demotion as relative to thresholds as in [9], or to be considered as trading-off against one another as in [40]?

3. Should values be considered separately as in [9], pairwise, as in [15], or collected together as in [29] and [24]?

Some suggestions were made in [12], but none of the possibilities as yet command general acceptance.

In general there are practical problems with factor reduction in that reducing the available factors results in a coarser grain representation which may not be adequate to capture significant differences between cases, and in consequence there may be difficulties in maintaining a consistent precedent base, given the reduced degrees of freedom. Moreover there is the theoretical problem of comparing sets of values: this does not arise for sets of factors since each preference (rules of the form of R3, as given in section 3.1) relates to two specific sets decided in a particular case, so general principles of comparison are not required.

3.3 Partitioning

The third approach is to decompose the problem into several partitions. Suppose we can decompose the problem into a small number of issues, and disjoint sets of factors relate to these issues. This will greatly reduce the potential comparisons. For example, as noted in [23] the Restatement of Torts states two requirements to find for the plaintiff: that the information was a trade secret and that it was misappropriated. Suppose that 6 of the 13 factors for the plaintiff and defendant relate to the first issue and the other 7 factors of each type to the second issue. Now there are 64 possible sets for each party relating to the first issue and 128 to the second, giving 2^{12} (4096) and 2^{14} (16384) potential comparisons respectively. Because we have a logical relation between the issues (both have to be decided for the plaintiff for the case to be decided for the plaintiff) we can consider them independently, so that we have their sum rather than their product to deal with. While 20480 is still a large number of comparisons, it is a significant improvement on the 67108864 required without partitioning. Moreover, we can improve on this by increasing the number of partitions: IBP [23] in fact uses five issues. Suppose each issue relates to three plaintiff and three defendant factors (some factors relate to more than one issue). Now we have only 64 comparisons for each issue - a total of only 320, less than twice the number of cases used in [23]. Now, especially when we see that some comparisons will be settled by *a fortiori* reasoning, that IBP makes only a single abstention in [23] becomes explicable. Note also the importance of treating the relationship between issues as a propositional function. If it mattered which factors were used to resolve the two issues so that we needed to balance a win for the plaintiff on one issue against a win for the defendant on the other, the number of comparisons would increase dramatically: in the two issue case it would require up to an additional 2^{14} (16384) comparisons to resolve situations in which one issue was found for each side.

The Abstract Dialectical Framework approach of [2] takes the partitioning approach further, by treating abstract factors as well as issues in this way. An ADF is defined in [22] as follows:

> **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.

It is helpful to partition L into supporting ($L+$) and attacking ($L-$) links. We can directly map the factor hierarchies of [4] into this structure, so that the set of nodes, S, comprises issues, abstract factors and base level factors, with base level factors forming the leaf nodes. This enables each abstract factor to be resolved separately. The ADF derived from CATO's factor hierarchy in [3] had 18 nodes defined in terms of base level factors. The number of children for the nodes varied: 1 had a single child, 9 had two children, 2 had three, 2 had four, 2 had six and 1 had seven. This partitioning gives a total of 330 potential comparisons. In fact this could be reduced by further decomposing some of the larger nodes (which we will explore in the next section).

In [3] all the nodes were considered expressible as propositional functions of their children: at this level of granularity, it becomes clear that precedents are not always needed to resolve comparisons: for example the abstract factor *EffortsToMaintainSecrecyVisAVisOutsiders* has one attacker (*SecretsDisclosedOutsiders*) and one supporter (*OutsiderDisclosuresRestricted*). It is obvious from the factor definitions that the abstract factor does not apply if its attacker is present unless its supporter is also present. That precedents would not always be required to resolve the comparisons was recognised in [4]:

> for certain conflicts, it is self evident how they should be resolved. ... It is not necessary to look to past cases to support that point. (p47).

It is possible to decompose the problem further. Any propositional expression can be expressed in Conjunctive Normal Form (CNF), the conjunction of a set of clauses in which each clause is disjunction of (positive or negative) literals. It is possible to rewrite any CNF formula so that every clause has exactly three literals, although there will typically be some cost in terms of additional literals and additional clauses. In a similar fashion we could introduce additional nodes to the ADF so that *every* non-leaf node has exactly two children. There are some advantages (with regard to manipulating and proving properties about the structure) in having a fixed number of children for every non-leaf node, so we introduce the notion of a *regular* ADF[9]. A *regular* ADF has the same number of children for every non-leaf node, and a *k-regular* ADF has exactly k children for every non-leaf node. Formally,

> **Definition 2**: For an $ADF\ A = < S, L, C >$, let L_{s_i} be the set of links to $s_i \in S$. A is a *regular ADF* if and only if the cardinality of L_s is the same for all $s \in S$ for which $L_s \neq \varnothing$. A is *k-regular* if and only if the cardinality of set L_s for every $s \in S$ is either 0 or k, and *2-regular* when $k = 2$.

We illustrate the reduction of the number of children using an example. In [4] and [2] the abstract factor *InfoUsed* has three supporters, (all of which are base level factors): F7 (*BroughtTools*), F18 (*IdenticalProducts*), and F8 (*CompetitiveAdvantage*), and one attacker, also a base level factor, F17, *InfoIndependentlyGenerated*. To give every node just two children, we introduce additional nodes *GaveHelp*, and *GaveAdvantage*. *InfoUsed* has *GaveHelp* as its supporter and F17 as an attacker. *GaveHelp* has two supporters, F8 and *GaveAdvantage* . Finally *GaveAdvantage* is supported by F7 and F18. In this way we have introduced two additional nodes, but reduced the number of comparisons. There may sometimes be an element of contrivance in the additional factors, but the computational gains compensate. Moreover, by reducing the number of base level factors at each node, we make it much more likely that there will be a precedent in the case base. The 2-regular ADF has 37 nodes, and there are a maximum of four different patterns to cover. In the next section we will produce a 2-Regular ADF for CATO, based on the analysis of [4].

4. 2-REGULAR ADF FOR CATO

CATO contains 26 base-level factors. These are shown in Table 2, together with the side favoured by each factor and the value

[9]The relationship with CNF may well be worthy of further exploration. Deciding whether a CNF formula with clauses of length 3 (3SAT) is satisfiable (there is an assignment of literals that makes every clause in the formula true) is a standard instance of an NP-complete problem [28], but there are effective algorithms for finding heuristic solutions. Several of these are based on the Davis-Putnam algorithm [26], and there are now a number of effective publicly available theorem provers which can be used to solve practical problems expressed as SAT problems. This might suggest an effective means of implementation.

which inclusion of the factor in the reasoning is considered to promote. These values are (largely[10]) taken from [25]. CA represents the value of respecting Confidentiality Agreements, MW, the value of protecting the Material Worth of the information to the plaintiff, LM, the value of allowing the defendant to use Legitimate Means to execute his business, QM, the value of discouraging the defendant from using Questionable Means, and RE the requirement that the plaintiff make Reasonable Efforts to protect the secret.

Table 2: Base Level Factors in CATO

ID	Factor	Value
F1	DisclosureInNegotiations (d)	RE
F2	BribeEmployee (p)	QM
F3	EmployeeSoleDeveloper (d)	LM
F4	AgreedNotToDisclose (p)	CA
F5	AgreementNotSpecific (d)	CA
F6	SecurityMeasures (p)	RE
F7	BroughtTools (p)	MW
F8	CompetitiveAdvantage (p)	MW
F9	Not used in [4]	
F10	SecretsDisclosedOutsiders (d)	RE
F11	VerticalKnowledge (d)	LM
F12	OutsiderDisclosuresRestricted (p)	CA
F13	NoncompetitionAgreement (p)	CA
F14	RestrictedMaterialsUsed (p)	QM/CA
F15	UniqueProduct (p)	MW/RE
F16	InfoReverseEngineerable (d)	LM
F17	InfoIndependentlyGenerated (d)	LM
F18	IdenticalProducts (p)	MW
F19	NoSecurityMeasures (d)	RE
F20	InfoKnownToCompetitors (d)	LM
F21	KnewInfoConfidential (p)	CA
F22	InvasiveTechniques (p)	QM
F23	WaiverOf Confidentiality (d)	CA
F24	InfoObtainableElsewhere (d)	LM
F25	InfoReverseEngineered (d)	LM
F26	Deception (p)	QM
F27	DisclosureInPublicForum (d)	LM/RE

We now construct a 2-regular ADF with these factors as leaf nodes. The nodes in this ADF will include nodes corresponding to all the abstract factors and issues given in [4] and all the issues in the logical model of [23], together with some additional nodes required to enable the regular structure (as discussed for *InfoUsed* towards the end of section 3). The complete set of nodes, children (and their associated values) is shown in Table 3.

4.1 Acceptance Conditions

If factors can only be present or absent, there are four possibilities for each node: both present, both absent, and the two ways in which one can be present and the other absent. Obvious ways of combining these are conjunction $(P \wedge Q) \rightarrow R$, for when both children are required and disjunction $(P \vee Q) \rightarrow R$ for when the children give independent ways of establishing the parent. In the other cases the children have opposite influences on the parent. Sometimes this is best seen as a way of representing an *exception*: i.e. R is generally false, but true when Q unless P. A paradigmatic example is when there is an agreement and a waiver: the waiver is the exception, nullifying the agreement and returning to the default situation. Different is the *rebuttal*: here there is no real default, R needs to be established by a factor. Q may offer a way of establishing R, but P can establish $\neg R$, and P may be preferred to Q. For example, independent development of the knowledge would rebut claims that it had been obtained by improper means. Both exceptions and rebuttals have the structure $(\neg P \rightarrow (Q \rightarrow R))$ so that the logic is the same (for Boolean factors) but their use in argumentation differs, for example with regard to burden of proof. If we use *unknown* as a truth value, rebuttals would have this as their default.

Suggestions for the appropriate functions for the nodes in the 2-regular ADF are included in Table 3. In most cases the choice between conjunction and disjunction is obvious: for example that bribery and burglary are alternative illegal acts (node 31), or that a valid agreement has to both have been made and still be in force (node 35). Choice between exception and rebuttal is more of a matter of judgement: since it does not affect the logic or the outcome, there is no obvious objective criterion, and for many purposes, such as outcome prediction, the difference is unimportant. The impact of the distinction comes in the way arguments are expressed, which requires detailed examination of particular decisions, and which matters only for the explanation of outputs.

If we allow the acceptance conditions to be expressed only as Boolean functions, the role of the precedents is simply to establish the nature of the function. For example, we should expect that in all cases where there was an agreement not to disclose (F4), it was the case that a confidentiality agreement existed, unless this had been waived (F23). Some factors, however, look suspiciously difficult to ascribe to cases in this binary fashion. Consider, for example, reverse engineerability (F16). Perhaps it would be possible, given enough time and expertise, to reverse engineer anything at all. But this is not what is intended by F16. Reading decisions of cases where this factor is present suggests that we need to think in terms of degrees of presence for such factors. Several decisions (e.g. *Televation*[11]) speak of the secret being "readily ascertainable". *Televation* concludes

> The mere fact, however, that a competitor could, through reverse engineering, duplicate plaintiff's product does not preclude a finding that plaintiff's techniques or schematics were trade secrets, particularly where, as here, the evidence demonstrated that the reverse engineering process would be time-consuming.

Moreover in *Mason*[12] the decision states (citing *KFC v Marion Kay*[13] and *Sperry Rand v Rothlein*[14] as support) that

> In this regard, we note that courts have protected information as a trade secret despite evidence that such information could be easily duplicated by others competent in the given field.

From this we have to think that the ascription of factors is not always a Boolean, black and white, affair, but rather may sometimes be a matter of degree. But this is of crucial importance: when we express the hierarchy of CATO as an ADF, the presence of F16

[10] F7 has been reassigned to MW, which seems to us to better reflect its purpose. Three others are associated with an additional value, to reflect the different ways they can be used to argue in different contexts.

[11] Televation Telecommunication Systems, Inc. v. Saindon, 522 N.E.2d 1359 (Ill.App. 2 Dist. 1988)

[12] Mason v. Jack Daniel Distillery, 518 So.2d 130 (Ala.Civ.App.1987)

[13] KFC Corp. v. MarionKay Co., Inc., 620 F.Supp 1160 (D.C.Ind 1985)

[14] Sperry Rand Corp. v. Rothlein, 241 F.Supp. 549 (D.Conn.1964)

Table 3: 2-Regular ADF for CATO. Nodes in italics correspond to issues in [23], nodes in bold face correspond to abstract factors in [4] and plain text nodes have been added to attain a 2-regular ADF.

ID	Node	Child 1	Child 2	Value 1	Value 2	Suggested Function
1	*TradeSecretMisappropriation*	SecretMisappropriated+	F3 -	RE + MW + (CA or QM)	LM	Rebuttal
2	SecretMisappropriated	Info Miasappropriated +	Info Trade Secret +	MW + (CA or QM)	RE + MW +	Conjunction
3	*Info Miasappropriated*	BreachOfConfidence +	ImproperMeans +	MW + CA	QM	Disjunction
4	*Info Trade Secret*	InfoValuable +	EffortstoMaintainSecrecy+	MW	RE	Conjunction
5	**InfoValuable**	InfoUseful +	KnownOrAvailable -	MW	LM	Exception
6	**EffortstoMaintainSecrecy**	AdequateEfforts+	SecurityFailures-	RE	RE	Rebuttal
7	InfoUseful	F8+	F15+	MW	MW	Disjunction
8	**KnownOrAvailable**	Known+	InfoAvailableElsewhere+	LM	LM	Disjunction
9	**Known**	KnownOutside+	Limitations-	LM	RE or CA	Exception
10	**InfoAvailableElsewhere**	F16+	F24+	LM	LM	Disjunction
11	KnownOutside	F20+	F27+	LM	LM	Disjunction
12	Limitations	F15+	MaintainSecrecyOutsiders+	RE	CA	Disjunction
13	**MaintainSecrecyOutsiders**	F10-	F12+	RE	CA	Exception
14	AdequateEfforts	F6+	MaintainSecrecyDefendant+	RE	RE	Disjunction
15	SecurityFailures	Reckless+	F21-	RE	CA	Exception
16	**MaintainSecrecyDefendant**	F4+	F1-	CA	RE	Exception
17	Reckless	F19+	Disclosed +	RE	RE	Disjunction
18	Disclosed	F27+	MaintainSecrecyOutsiders-	RE	RE	Exception
19	BreachOfConfidence	InfoUsed+	ConfidentialRelationship+	MW	CA	Conjunction
20	**ImproperMeans**	QuestionableMeans +	LegitimatelyObtainable-	QM	LM	Rebuttal
21	**ConfidentialRelationship**	NoticeofConfidentiality+	ConfidentialityAgreement+	CA	CA	Disjunction
22	**ConfidentialityAgreement**	F4+	F23-	CA	CA	Exception
23	**InfoUsed**	GaveHelp +	F17-	MW	LM	Rebuttal
24	**NoticeofConfidentiality**	ValidAgreement+	AwareConfidential +	CA	CA	Disjunction
25	**QuestionableMeans**	IllegalMethods+	DefOKMethods-	QM	RE or LM	Rebuttal
26	**LegitimatelyObtainable**	InfoKnownorAvailable+	InfoAvailableElsewhere+	LM	LM	Disjunction
27	GaveHelp	F18+	GaveAdvantage+	MW	MW	Disjunction
28	GaveAdvantage	F7+	F8+	MW	MW	Disjunction
29	IllegalMethods	Criminal +	Dubious +	QM	QM	Disjunction
30	DefOKMethods	F1+	DefendantDiscovered+	RE	LM	Disjunction
31	Criminal	F2+	F22+	QM	QM	Disjunction
32	Dubious	F14+	F26+	QM	QM	Disjunction
33	DefendantDiscovered	F17+	F25+	LM	LM	Disjunction
34	AwareConfidential	F14+	F21+	CA	CA	Disjunction
35	ValidAgreement	AgreementMade+	AgreementInForce+	CA	CA	Conjunction
36	AgreementMade	F4+	F13+	CA	CA	Conjunction
37	AgreementInForce	F5-	F23-	CA	CA	Conjunction

appears to be immediately decisive for the defendant (and yet a number of cases with F16 were, in practice, found for the plaintiff). Now if we think of the possibility of factors being present to varying degrees, we get further possibilities, including (for conjunction or disjunction):

- We might need both children present to a sufficient degree, where these degrees are independent;
- We might need both children present, but a greater degree for one could compensate for a lesser degree of the other (i.e. we can sum them, perhaps even as a weighted sum);
- We might need either child present to a certain degree;
- We might need the combination to reach a certain degree, but this could come from either factor or a combination of the two.

For exceptions and rebuttals, as well as the value to which the factors relate, we might want to consider:

- The exception (or rebuttal) might need to attain a sufficient degree;
- The degree the exception (rebuttal) needs to attain might depend on the degree of the other factor;
- Since rebuttal expresses a preference, we might need to compare the degrees of the factors involved.

For the present we must leave this for future work: both the question of degrees of presence and how this needs to be treated will require some detailed consideration of decisions and empirical evaluation to determine the effect of different treatments. The questions are reminiscent of the questions relating to values posed in [12], questions of accrual as addressed in [36] and [1], and ways in which the ascription of factors can be argued for [38], in which HYPO's notion of dimensions is revived to be used as a link between facts and factors. Note, however, that we are dealing with (at most) two elements, and so at least the difficulties associated with comparisons between sets do not arise.

5. WHERE ARE THE VALUES?

We are now in a position to return to values. In [25] factors were associated with five values. The rationale given there for these values was that Trade Secret Law existed in order to encourage innovation, and so needed to protect the investment of time and effort in new techniques and products. But it is undesirable that there be too much litigation: thus the putative secrets needed to have some *Material Worth* (MW). Moreover, again to minimise litigation, the developer could be expected to adopt some (reasonably) adequate security measures, that is to make *Reasonable Efforts* (RE) to protect the secret. Ideally this could take the form of explicit *Confidentiality Agreements* (CA) which should make the situation clear to all concerned and the law should enforce these agreements where they exist. The law should, however, also protect the developer against dubious activity and the use of illegal or *Questionable Means* (QM) should not be countenanced. On the other hand it is not the place of the law to discourage enterprise and exploitation of innovative ideas, and so if the information was acquired by *Legitimate Means* (LM), this should be a defence. An attribution of values to factors is shown in Table 3: this follows [25] with the exceptions noted in footnote 9.

5.1 Adding Values to the ADF

The values discussed above can now be added to the ADF in Table 3. The leaf nodes are all base level factors, and so are assigned the values as in Table 2. The values of the children then determine the values of their parents as follows:

- Where both children relate to the same value, this value is assigned to the parent.
- For rebuttals and exceptions, it is the value of the child which leads to the *acceptance* of the adult which is propagated: where the rebuttal/exception holds, the parent is not present and so does not pass a value up the tree.
- For disjunctions, the value of the parent depends on which of the disjuncts is used.
- For conjunctions the parent receives the values of both children.

This helps us to understand the idea in [5] that values can be associated with rules *and* their individual conditions. The acceptance conditions of each non-leaf node can be seen as one or more rules. The value promoted by these rules is the value assigned to the node. Note that this does mean that the value of the rules may be contextually dependent when the node has disjunctive acceptance conditions, and the children have different values. This occurs three times in our ADF for CATO (nodes 3, 9 and 30). For node 3 there are two distinct ways in which misappropriation can be established - either by breaching a confidence (assuming that the information has some material worth), or by some illegal act. We want the Trade Secrets Law to promote both the values of CA (given MW) and QM, but the cases form two natural partitions according to the nature of the wrongdoing alleged. For node 9 the effect of the information being known to outsiders can be shown limited by the plaintiff making reasonable efforts as demonstrated by the uniqueness of the product, or by an explicit confidentiality agreement. Finally, there are two different ways in which the actions of the defendant can be vindicated in node 30: either the plaintiff was remiss (RE), or the defendant discovered the information legitimately (LM). The disjunctive node 30 allows either value to be promoted, as fits the facts. Where conjunctions have children with different values, part of the purpose of the rule is to ensure that both are considered. Thus for node 1, the top level of the rule relates to all five values (although CA and QM are alternatives for the reasons just discussed). These conjunctions with different values occur towards the top of the tree as the various strands of the argument are brought together.

The values of the children are the values promoted by the inclusion of the condition in the rule. This is of particular importance when the values differ and we have an exception or a rebuttal. There are nine such nodes. Consider, for example, nodes 20 and 25. Here we want to promote QM by disallowing illegal methods, but to ensure that the values of RE and LM are properly accounted for: hence the possibility of rebuttal based on these values. In node 23 we want to pay proper attention to MW and so require that the information was of assistance to the defendant, but again we need to ensure that the use of LM can vindicate the defendant. In node 16 we recognise that disclosure would suggest negligence on the part of the plaintiff, but since disclosure in negotiations may be essential in order to do business, we want to allow the plaintiff the means to forestall the charge of negligence. In nodes 15 and 13 it is recognised that a duty of "fidelity and trust" may exist, even when the plaintiff has been negligent (this often applies when the defendant was given the information while employed by the plaintiff as in *Space Aero*[15]). Node 9 is primarily intended to advance LM, but again recognises that the plaintiff may protect a secret even when it is known to others. Node 5 recognises, however, that the worth of a secret may be lost when the information is known, and these steps have not been taken. Finally node 1 recognises that even when all else is in order for the plaintiff, there may still be a defence based on the fact that the defendant did nothing wrong. Note that in these cases we need the rule to enable the law to promote value of the main condition, but we need the exception so that some other value can be considered. Note too that the recognition of these nodes as rebuttals and exceptions commits to a preference for the value of the rebuttal/exception. If we wanted to place more stress on the need for plaintiffs to protect their own secrets, we would not allow the exception in node 15, but demand that plaintiffs require their employees to sign formal non-disclosure agreements. (If such an agreement had existed in *SpaceAero*, node 17 would not be accepted because the non-competition agreement showed security measures had been taken and there was no public disclosure). Where we have rebuttals and exceptions in the same value, there is no need for preference or choice: for example in node 22, it is obvious that for there to be a valid confidentiality agreement it is necessary that the agreement has been made and not waived.

5.2 ADF as a Theory of the Case Law

So let us return to those nodes where there is a conflict of values. The discussion above suggests that these are the points in a decision where a tension between values is resolved: a balance is struck, or a preference is expressed. Moreover, the existence of the node recognises a preference: if the preference was not held, the exception or rebuttal would not be effective, and the node would not be required. What this suggests is that the ADF expresses a theory, already present in the case law, rather than requiring the construc-

[15] Space Aero Products Co. v. R.E. Darling Co., 238 Md. 93, 208 A.2d 74 (1965)

tion of a theory, as suggested by [35] and [15]. This should not be seen as a problem. According to Levi [32], case law goes through a life cycle of three stages. First there is a period of flux, in which there are not enough precedents to provide guidance to the judges. Here the judges are genuinely engaged in theory construction, introducing factors, striking balances and expressing preferences. At this stage many of the cases can be seen as test cases, and many Supreme Court decisions are part of this stage. But as time goes on, the balances and preferences become entrenched through the decisions taken and the theory becomes established, and almost always capable of resolving new cases. In this second stage there is a period of stability as the theory is applied. Eventually, however, perhaps because of a change in society, or its attitudes, the theory becomes increasingly challenged and the stability breaks down. See [17] for a discussion of the signs that this third stage in being entered upon. At some point the tensions become too great for the theory to be sustained and the first stage of a new theory begins. The reasoning expressed by CATO and formalised in [30] is, as [30] makes clear, intended only to apply to the second stage: that is to encapsulate and apply a settled doctrine. To be in this stage, we should be able to cite a precedent justifying our choice of acceptance conditions at each node. But this has an upper bound of 4*37 = 148, since (given factors either present or absent), each node has a maximum of four relevant cases. This further suggests that the size of the precedent base in CATO and IBP was about right[16] (assuming the precedents included to be well chosen).

5.3 Challenging the Theory

The theory, expressed as the ADF, does indicate, however, points at which we can challenge the theory. This depends on the nature of the node. Cases include:

- A conjunction or disjunction can be challenged by finding an exception or rebuttal. We may well think that node 10 is too favorable to the defendant, and that the plaintiff should have some way of defending against this. For example, there could be an exception if questionable methods were used in the reverse engineering. It may be that the information was only reverse engineerable with the assistance of restricted materials (e.g *Mineral Deposits*[17], in which after one of the defendants "received the spiral, he removed the label which indicated that patent applications were pending and gave the spiral to defendant Zbikowski. Zbikowski then cut the spiral into pieces, made molds of the components, and proceeded to manufacture copies of the spiral."). We may therefore wish to express the importance of QM by allowing F14 as an exception to F16 here.
- We might want to add an additional disjunct. For example the use of blackmail in a new case would suggest that we need another alternative at node 31.
- We may wish to claim that the value preference no longer reflects the attitudes of society (as discussed in [12]). It may be that we feel node 15 is too favourable to the plaintiff, inhibiting competition, so that we do not want to allow F21 as an exception, effectively preferring LM to CA at this node.

5.4 Example

For an example we will use *The Boeing Company v. Sierracin Corporation*, which involved cockpit windows. Factors present, according to [4] are F4 (AgreedNotToDisclose), F6 (SecurityMeasures), F12 (OutsiderDisclosuresRestricted), F14 (RestrictedMaterialsUsed) and F21 (KnewInfoConfidential) for the plaintiff and F1 (DisclosureInNegotiations) and F10 (DisclosureInNegotiations) for the defendant. The information is accepted as useful (node 7), satisfying MW, and that secrecy had been maintained with respect to outsiders (node 13) is established since any disclosures (F1 and F10) had been covered by a non-disclosure agreement (F4), giving preference to CA over RE. There is, therefore a trade secret, adequate efforts to maintain secrecy having been taken (F6 at node 14), so that RE is also respected. Misappropriation is based on a breach of confidentiality (node 19), since it was accepted that the information had been used, and there was a confidential relationship (node 21), as established by the non-disclosure agreement (F4 at node 22) In the absence of an explicit agreement, confidentiality could have been established using F14 or F21 node 34, but then the case may have failed at node 13 because of the disclosures to outsiders in F10. In the actual judgement the defendants were given an opportunity to reverse engineer the product during a 120 day grace period, but failed to do so, preventing any argument based on LM. Finding that questionable means were used is unnecessary, although arguable using F14 at node 32. Thus the finding for the plaintiff promoted RE, CA, MW.

5.5 Degrees of Presence

All of this assumes that the presence or absence of a factor is Boolean: if not, the ability to give definitive answers evaporates. But as discussed in section 4, there are excellent reasons to see some factors as present to different degrees. Now we have reconsidered the role of values, we may wish to express these differing degrees in terms of values. The degree to which the presence of a factor promotes a value is dependent on the degree to which the factor is present. Thus, consider F16, reverse engineerable. If it is obvious from inspection how to copy make the project (i.e. F16 is fully present) then LM is strongly promoted. Conversely, if reverse engineering would involve a great deal of highly skilled effort (i.e. F16 is present only to a very limited degree), LM would be promoted only to a very limited extent by accepting node 10. The advantage of moving from factors to values in this way is that we can also consider the degree to which factors promote their values. For example, both F13 (that there is a non-competition agreement) and F21 (that the defendant was aware that the information was confidential) promote CA, but it is clear that F13 promotes CA to a greater degree than F21. We can now see the degree to which a value is promoted as a function of the specific factor promoting it (e.g. F24 will always promote LM to a greater extent than F16) *and* the degree to which the factor is present. For example whether LM is promoted more by F20 or F27 depends on the number of competitors to which the information is known and the accessibility of the public forum in which it was disclosed. The questions of section 4 remain, but now they can be posed in terms of values rather than factors. Importantly, whatever the precise mechanism chosen, this allows for an exception to exist in general, but be capable of rejection after consideration because its acceptance would insufficiently promote the value[18].

[16] We assume that the fact that CATO contained 148 cases was a coincidence, since it is not justified in this way in [4].

[17] Mineral Deposits Ltd. v. Zigan, 773 P.2d 606 (Colo.App.1988)

[18] Consider for example the holding in *Avecedo v California* (500 U.S. 565 (1991)) that "The doctrine of *stare decisi*s does not preclude this Court from eliminating the warrant requirement of Sanders, which was specifically undermined in Ross. The Chadwick-Sanders rule affords *minimal protection to privacy interests*" (emphasis ours). This is a very clear acknowledgment that values can be promoted to different degrees, and that an exception applies only when they are sufficiently promoted.

While the degree of promotion (or at least the *relative* degree of promotion) can be determined independently of the facts which led to the ascription of the factor to the case (see [25] for one attempt to do this in the CATO domain), we cannot determine the degree of presence of a factor without a close examination of the facts. For this reason it is important that we are able to challenge the ascription of factors on the basis of the particular facts of the case. In fact, the need for this has already been recognised: as [11] observes, the famous case of *Pierson v Post* (which was used in [16] to motivate the whole line of work on purposes and values) turns on the ascription of factors rather than their comparison. Approaches to exposing the ascription of factors to argumentation have been suggested in [10] and [2].

Note also that there are several points in the theory at which a balance needs to be struck, several of them quite low down in the tree. At the top level, rules tend to express the need to accommodate several values in the reasoning (thus in Table 3, node 1 refers to all five values). This accords both with work such as [15] in which the value conflicts on which the case turns do indeed appear low in the tree, and [31], which observes that "often balances are encountered at particular nodes in a process that involves rules and other reasoning strategies". Moreover, the account presented here allows for different preferences to be expressed at different nodes, so that assessment of the relative importance of pair of values can be made contextually dependent. This also serves to remind us of the arguments in favour of reasoning with portions of precedents [19], an approach which has been somewhat neglected in recent work.

6. CONCLUDING REMARKS

In this paper we have reconsidered the topic of reasoning with legal cases represented as sets of factors, with the particular perspective of how it is possible to go beyond what is already present in the precedents. We considered three main approaches: simple rule broadening, and two methods of reducing the comparisons, (which are exponential in the number of factors). Comparisons can be reduced by a reduction in the number of factors, and by partitioning, so that the problem is decomposed into smaller problems, each with fewer factors, although the total number of factors remains the same. Our view was that rule broadening often lacked justification, that reduction in the number of factors could result in too great a loss of expressiveness and hence the inability to capture relevant distinctions. We therefore adopted a partitioning strategy, and ensured that our partitions were as small as possible by making our chosen representation a 2-regular ADF. In this way we never need to compare sets of factors or values, only pairs. We applied this approach to the well known and documented CATO system [4], to act as a source of concrete illustrations. Once we had developed this structure, we were able to be more precise about the role of values and purposes in the reasoning: where they entered the picture, at which points they were deployed, and the way in which preferences could depend on context.

We also noted that the success of systems based on this approach required both

- that the domain modelled be in a stable state (the second stage of Levi's life cycle explained in [32]);
- that it is possible to assign factors to cases in a Boolean manner, without varying degrees of presence.

For the future we believe that it is the second of these requirements that needs the more careful attention, since the first is often satisfied for a considerable period of time. Two things are required for the second problem:

1. incorporation of a layer of facts in the representation, so that cases can be represented in terms of their facts, rather than as sets of factors, so that different degrees of presence can be explicitly challenged and justified.
2. a principled approach to the questions in section 4, so that the effect of different degrees of presence can be handled. When considering values, we noted that values could also be promoted to differing degrees, but that since differing degrees of presence for factors could be captured by different degrees of promotion of values, factors and values could be considered together with respect to this issue.

Note, however, that the role of precedents on this account is rather static. The theory is constructed by the analyst when constructing the ADF, not by the system at run time, as was the aspiration of [24]. Adding a new precedent to the precedent base will not cause the system to adapt to this decision. Precedents can be used to justify what the theory contains, but the theory is fixed. Therefore, if we have a case which diverges from the decision predicted by our ADF which we need to accommodate, we need to modify the ADF in some way (see [3]). To do this it would be helpful to be able to argue about preferences (perhaps using some form of meta-level argumentation as in [14]).

Levi argues that legal reasoning in a particular domain has a life cycle in which a period of stability eventually breaks down. Perhaps reasoning about legal reasoning has a similar life cycle. For nearly two decades many researchers have been content to view cases as represented in terms of a sets of factors, and to explore the logic of resolving conflicts and drawing legal consequences from them. This has yielded many insights, but is now seen to have limitations. The need to acknowledge different degrees (of factor presence or of value promotion), the need to argue that value preferences require contextualisation [29] and modification, and the need to argue in terms of facts as well as factors, have all applied pressures that require cases to be seen as more than sets of factors. In this paper we have attempted to come up with a framework which can act as a possible setting within which these issues can be explored further.

References

[1] L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. A strategy for deliberation exploiting accrual of values. In *Proceedings of the Ninth International Workshop on Argumentation in Multi-Agent Systems (ARGMAS 2012)*, pages 1–19, 2012.

[2] L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Abstract dialectical frameworks for legal reasoning. In *Proceedings of JURIX 2014*, pages 61–70. IOS Press, 2014.

[3] L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Evaluating the use of abstract dialectical frameworks to represent case law. In *Proceedings of the 15th ICAIL*. This volume, 2015.

[4] V. Aleven. *Teaching case-based argumentation through a model and examples*. PhD thesis, Univ.of Pittsburgh, 1997.

[5] M. Araszkiewicz, A. Lopatkiewicz, and A. Zienkiewicz. Factor-based parent plan support system. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pages 171–175, 2013.

[6] K. Ashley. *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. MIT Press, Cambridge, MA, 1990.

[7] K. Ashley. Teaching a process model of legal argument with hypotheticals. *Artif. Intell. Law*, 17(4):321–370, 2009.

[8] K. D. Ashley and S. Brüninghaus. A predictive role for intermediate legal concepts. In *Proceedings of JURIX 2003*, pages 153–162. IOS Press, 2003.

[9] T. Bench-Capon and H. Prakken. Using argument schemes for hypothetical reasoning in law. *Artif. Intell. Law*, 18(2):153–174, 2010.

[10] 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.

[11] T. Bench-Capon and E. Rissland. Back to the future: dimensions revisited. In *Proceedings of JURIX 2001*, pages 41–52. IOS Press, 2001.

[12] T. J. M. Bench-Capon. Relating values in a series of supreme court decisions. In *Legal Knowledge and Information Systems - JURIX 2011: The Twenty-Fourth Annual Conference, University of Vienna, Austria, 14th-16th December 2011*, pages 13–22, 2011.

[13] T. J. M. Bench-Capon, K. Atkinson, and A. Chorley. Persuasion and value in legal argument. *J. Log. Comput.*, 15(6):1075–1097, 2005.

[14] T. J. M. Bench-Capon and S. Modgil. Case law in extended argumentation frameworks. In *The 12th International Conference on Artificial Intelligence and Law, Proceedings of the Conference, June 8-12, 2009, Barcelona, Spain*, pages 118–127. ACM Press, 2009.

[15] T. J. M. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artif. Intell.*, 150(1-2):97–143, 2003.

[16] D. Berman and C. Hafner. Representing teleological structure in case-based legal reasoning: The missing link. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pages 50–59. ACM Press, 1993.

[17] D. H. Berman and C. D. Hafner. Understanding precedents in a temporal context of evolving legal doctrine. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 42–51. ACM Press, 1995.

[18] F. J. Bex. *Arguments, stories and criminal evidence: A formal hybrid theory*, volume 92. Springer, 2011.

[19] L. K. Branting. Reasoning with portions of precedents. In *Proceedings of the 3rd ICAIL*, pages 145–154. ACM, 1991.

[20] L. K. Branting. A reduction-graph model of ratio decidendi. In *Proceedings of the Fourth International Conference on Artificial intelligence and Law*, pages 40–49, 1993.

[21] J. Breuker and N. den Haan. Separating world and regulation knowledge: Where is the logic. In *Proceedings of the Third International Conference on Artificial Intelligence and Law*, pages 92–97. ACM Press, 1991.

[22] G. Brewka, H. Strass, S. Ellmauthaler, J. Wallner, and S. Woltran. Abstract dialectical frameworks revisited. In *Proceedings of IJCAI 2013*, 2013. Issued on CD.

[23] S. Brüninghaus and K. Ashley. Predicting outcomes of case-based legal arguments. In *Proceedings of the Ninth International Conference on Artificial Intelligence and Law*, pages 233–242. ACM Press, 2003.

[24] A. Chorley and T. Bench-Capon. Agatha: Using heuristic search to automate the construction of case law theories. *Artif. Intell. Law*, 13(1):9–51, 2005.

[25] A. Chorley and T. Bench-Capon. An empirical investigation of reasoning with legal cases through theory construction and application. *Artif. Intell. Law*, 13(3-4):323–371, 2005.

[26] M. Davis and H. Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, July 1960.

[27] A. v. d. L. Gardner. *Artificial intelligence approach to legal reasoning*. MIT Press, Cambridge, MA, 1984.

[28] M. R. Garey and D. S. Johnson. *Computers and intractability*. W.H. Freeman, 1979.

[29] M. Grabmair and K. D. Ashley. Facilitating case comparison using value judgments and intermediate legal concepts. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 161–170. ACM Press, 2011.

[30] J. Horty and T. Bench-Capon. A factor-based definition of precedential constraint. *Artif. Intell. Law*, 20(2):181–214, 2012.

[31] M. Lauritsen. On balance. *Artif. Intell. Law*, 23(1), 2015.

[32] E. H. Levi. *An introduction to legal reasoning*. University of Chicago Press, 1949.

[33] L. Lindahl and J. Odelstad. Intermediaries and intervenients in normative systems. *J. Applied Logic*, 6(2):229–250, 2008.

[34] L. T. McCarty. Reflections on taxman: An experiment in artificial intelligence and legal reasoning. *Harvard Law Review*, pages 837–893, 1977.

[35] L. T. McCarty. An implementation of Eisner v. Macomber. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 276–286. ACM Press, 1995.

[36] H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the 10th ICAIL*, pages 85–94. ACM Press, 2005.

[37] H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artif. Intell. Law*, 6(2-4):231–287, 1998.

[38] 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*, 2013, Available online.

[39] E. L. Rissland and K. D. Ashley. A note on dimensions and factors. *Artif. Intell. Law*, 10(1-3):65–77, 2002.

[40] G. Sartor. Doing justice to rights and values: teleological reasoning and proportionality. *Artif. Intell. Law*, 18(2):175–215, 2010.

[41] D. B. Skalak and E. L. Rissland. Arguments and cases: An inevitable intertwining. *Artif. Intell. Law*, 1(1):3–44, 1992.