

Using Issues to Explain Legal Decisions

Trevor Bench-Capon

University of Liverpool, Liverpool, UK. tbc@csc.liv.ac.uk

Abstract. The need to explain the output from Machine Learning systems designed to predict the outcomes of legal cases has led to a renewed interest in the explanations offered by traditional AI and Law systems, especially those using factor based reasoning and precedent cases. In this paper we consider what sort of explanations we should expect from such systems, with a particular focus on the structure that can be provided by the use of issues in cases.

Keywords: Reasoning with cases · Explanation · Legal Reasoning · Factors · Issues

1 Introduction

Currently there is much interest in the use of Machine Learning (ML) based approaches to predict legal decisions. The European Convention on Human Rights alone has been the subject of a cluster of such systems including [2], [37], [20], [38] and [34]. Using such systems in legal applications, however, raises a number of issues [9], including bias, adapting to changes in statute law, case law and social values, and, perhaps most important, the lack of explanation. In law there is a right to explanation [22] and so providing explanations that users can understand is essential if AI systems are to be used in practice. This has long been recognised in AI and Law and the provision of explanations has been a central feature of systems developed in the field [6]. It is therefore a natural move to see whether it is possible to use the techniques developed to explain the outputs of systems previously developed in AI and Law to explain the outputs of ML systems. In particular the factor based reasoning developed for the CATO system [3] and widely adopted in subsequent systems [7] has been proposed as a suitable candidate for this role. In [17] the idea is to first ascribe factors to cases using ML and then to explain the outcomes in terms of these factors. In [42], the proposal is to produce the explanation from a case based system running in parallel with the ML system.

In this paper we consider what sort of explanations we can expect from such systems, with a particular focus on the structure that can be provided by the use of issues in cases.

2 Background: Factor Based Reasoning in CATO

We begin by describing factor based reasoning in CATO [3], which is the starting point for subsequent accounts of factor based reasoning [7]. CATO is directed

Table 1. Factors in CATO. Numbers are as in [3]; “p” and “d” indicates the side favoured as in [28]

Plaintiff Factors		Defendant Factors	
F2p	Bribe-Employee	F1d	Disclosure-In-Negotiations
F4p	Agreed-Not-To-Disclose	F3d	Employee-Sole-Developer
F6p	Security-Measures	F5d	Agreement-not-specific
F7p	Brought-Tools	F10d	Secrets-Disclosed-Outsiders
F8p	Competitive-Advantage	F11d	Vertical-Knowledge
F12p	Outsider-Disclosures-Restricted	F16d	Info-Reverse-Engineerable
F13p	Noncompetition-Agreement	F17d	Info-Independently-Generated
F14p	Restricted-Materials-Used	F19d	No-Security-Measures
F15p	Unique-Product	F20d	Info-Known-to-Competitors
F18p	Identical-Products	F23d	Waiver-of-Confidentiality
F21p	Knew-Info-Confidential	F24d	Info-Obtainable-Elsewhere
F22p	Invasive-Techniques	F25d	Info-Reverse-Engineered
F26p	Deception	F27d	Disclosure-In-Public-Forum

towards the domain of US Trade Secrets, and is largely based on the law as set out in the *Restatement of Torts*¹. In CATO cases are represented as sets of factors. Factors are ascribed on the basis of stereotypical patterns of facts which have legal significance in that they provide a reason to find for one of the parties. CATO has thirteen factors for each side, as shown in Table 1. The conditions for ascribing them to cases are given in Appendix 2 of [3]. Note that the absence of a factor is not in general a reason to find for the other party: in the rare cases where the absence of a factor might favour the other side, a second distinct factor for the party favoured is used. The only example of this in CATO is *security measures* with factors F6p and F19d. Note, however, that it may be that neither F6p or F19d is present in the case: even if security measures were taken, so that there is no reason to find for the defendant on this aspect, they may not have been sufficient to provide a reason to find for the plaintiff, and so that the aspect is neutral. The factors from [3] have been reused by many subsequent researchers, including [19], [21], [1], [55] and [42].

When presented with a new case a three ply argument is constructed. In the first ply a proponent cites the most-on-point precedent (i.e. the precedent with the greatest overlap of factors irrespective of which side they favour) decided for the side being argued for. Suppose this is the plaintiff. In the second ply the opponent either cites a counterexample (a case which favours the defendant and is at least as on point as the case cited by the plaintiff) or distinguishes the precedent by pointing to a factor favouring the plaintiff in the precedent but not the new case, or a factor favouring the defendant in the current case but not the precedent. In the third ply the plaintiff offers a rebuttal by distinguishing the counterexamples, or downplaying the distinguishing factor by pointing to a

¹ The relevant section, section 757, *Liability for disclosure or use of another’s Trade Secret*, can be found at <https://www.lrdc.pitt.edu/ashley/restatem.htm>.

factor which can cancel the additional factor or a factor which can be substituted for the absent factor [44].

CATO organised its factors into a *factor hierarchy*. At the upper level are *issues*, and below these are layers of abstract factors before the leaf nodes are reached. These leaf nodes are the base level factors shown in Table 1. The importance of this hierarchy is for determining whether distinctions can be downplayed: a factor can only be substituted for or cancel another factor if they have a common ancestor. The closer the ancestor the more persuasive the downplay. How persuasive the downplay is matters for prediction, where the success or otherwise of the rebuttal needs to be decided, but not for explanation. The success or otherwise of the rebuttal is given by the outcome of the case, which shows whether or not the downplay was successful.

CATO’s factor based model has inspired a number of formal accounts of precedential constraint [29], [33], [45] and [41]. These models are based on a way of representing precedents originating in [43]. Suppose we have a case with plaintiff factors P and defendant factors D . Now the strongest reason to find for the plaintiff will be the conjunction of all the factors in P and the strongest reason to find for the defendant the conjunction of all the factors in D . The outcome of the case will show which reason was preferred. A decision is taken to be constrained in these approaches if deciding for the other party would introduce an inconsistency into the set of preferences in the precedent base². Using all the factors available for the winning side is termed the *results* model in [29]. It may be, however, that a subset of the factors for the winning side is still sufficient to overcome the reason for the losing side. This would allow a subset of the winner’s factors to be used in the preference. This is termed the *reason* model in [29]. A comparison of the two models is given in [41].

The reasoning in CATO: citation, distinguishing and counterexample, followed by rebuttal through downplaying distinctions and distinguishing counterexamples was expressed as a set of argumentation schemes in [54] and formalised in ASPIC+ in [44]. This formalisation uses the results model, and uses the full set of factors available to both sides. These schemes were proposed as a means of providing explanation for ML systems in [42]. We will discuss the explanations from [42] in the next section.

3 Explanation with Argument Schemes

Explanation in [42] takes the form of a dialogue modelling the three ply argumentation structure of CATO. An example dialogue is shown in Figure 1. Figure 1 illustrates a particular example used in [42]. The two precedent cases and the new case in that example are shown in Table 2. The cases have been given mnemonic names. The top layer, put forward by the proponent, is an argument based on citing a precedent case. The second layer, objections by the opponent, comprises objections based on each of the two types of distinction ($O1a$ and

² In practice this formal notion of constraint may not actually be respected in a given judicial setting. For a jurisprudential discussion see [50].




Fig. 1. Example Dialogue Tree From [42]

$O1c$), and a counter example ($O1b$). The final layer shows the proponent’s rebuttals: each objection is met by both a substitution and a cancellation ($P2a$ and $P2'a$ for $O1a$ and $P2c$ and $P2'c$ for $O1c$), and a rebuttal of the counterexample through a “transformation”, which means that substitutions and cancellations can transform the case into a precedent for the proponent’s side.

Table 2. Cases in the Example From [42]

Case	Outcome	Plaintiff Factors	Defendant Factors
Deceived	P	F6p F26p	F10d F24d
NoMeasures	D	F2p	F10d F24d
Bribed	TBA	F2p F6p	F10d F16d

The diagram in Figure 1 offers an explanation for a plaintiff win in *Bribe*. *Deceived* matches because the plaintiff took security measures (F6p) and disclosed information to outsiders (F10d). The defendant can now cite distinctions of both kinds: deception (F26p) was not used, and the information is re-engineerable (F16d) in the new case but not the precedent. Moreover *NoMeasures* also matches on two factors and so is as on point as *Deceived* and so can serve as a counter example. To counter $O1a$ the plaintiff now argues ($P2a$) that the lack of deception does not matter because bribery was used (F2p) and this can substitute for F26p. Alternatively it can be argued ($P2'a$) that the additional defendant factor in the precedent, F24d, that the information was available elsewhere, cancels the additional plaintiff strength of the precedent coming from bribery. It is clear in this case that the substitution is more effective: bribery and deception play similar roles, both being different examples of the use of improper means. In the case of the additional defendant factor in the new case,

F16d, used in *O1c*, it can be substituted by the additional factor in the precedent, F24d (*P2b*), or cancelled by the additional plaintiff factor, F2p (*P2'b*). Again it seems that substitution is the better argument because of the similarity of the roles of F16d and F24d.

The fact that the substitutions are clearly better rebuttals than the cancellations (in this case: for other examples the reverse will be true) highlights a problem with the explanation. Although the arguments generated by the definitions are possible arguments in that they conform to the definitions of [42], they lack plausibility because they relate to entirely different concerns. That is, these objections fail to make sense in domain terms. In CATO the strength of a downplay depended on how close the factors were in the factor hierarchy. In Figure 1, however, it is not possible to tell which rebuttal succeeded, or whether a combination of the two was required. All the candidate arguments are presented, but it is the user that must supply the domain knowledge to assess how strong these arguments are, and which objections and rebuttals should succeed in the particular case. The reason for the decision is there, but the user must extract it. In order to guide the user, we turn to consider the knowledge of the structure of the domain represented by the use of issues.

4 Issues

In CATO the top level of the factor hierarchy was made up of *issues*. The role of issues in a case is to identify the salient points that need to be shown in order to prove or defend a case, and, hence, the factors that are relevant to the different points. Issues in CATO served mainly to organise the explanation. CATO identified two main issues for Trade Secret misappropriation: the information had to be a trade secret (with the burden of proof on the defendant to show that it was not) and the information had to have been misappropriated (with the burden of proof on the plaintiff). Below these main issues were sub-issues. To be a trade secret the information had to be valuable and its secrecy adequately maintained. If the information was used, misappropriation could be shown either through a breach of confidence or through the use of improper means by the defendant. For a breach of confidence a confidential relation between plaintiff and defendant had to exist. Issues took on an additional significance when CATO was adapted to predict outcomes in the IBP system [19] when the issues formed a top layer of necessary and sufficient conditions (termed the logical model in [19]), with factor based reasoning used to determine the status of the leaf issues. This structure, strict logical rules at the top with case based reasoning to determine which rules applied, was earlier used in CABARET [51], and was later adopted and adapted to accommodate his value judgement formalism by Grabmair [27]. Figure 2 shows top level logical model and the allocation of factor to issues in [27]³.

³ We use structure of the logical model in [27] which differs slightly from that of [19].




Fig. 2. Issues and Factors From [27]. Factors may relate to more than one issue.

We can now consider the example cases in Table 2 in terms of issues. We also include in Table 3 some other actual cases we will mention in this paper⁴.

⁴ Mason v. Jack Daniel Distillery, 518 So.2d 130 (Ala.Civ.App.1987), Leo Silfen, Inc. v. Cream, 29 N.Y.2d 387, Computer Print Systems v. Lewis, 422 A.2d 148 (1980), K & G Oil Tool & Service Co. v. G & G Fishing Tool Serv., 314 S.W.2d 782, (1958), College Watercolor Group, Inc. v. William H. Newbauer, Inc., 468 Pa. 103, 360 A.2d 200 (1976), Arco Industries Corp. v. Chemcast Corp., 633 F.2d 435, 208 USPQ 190 (6th Cir.1980), E. V. Prentice Dryer Co. v. Northwest Dryer & Machinery Co., 246 Or. 78, 424 P.2d 227 (1967), Kinnear-Weed Corp. v. Humble Oil Refining Co. 150 F. Supp. 143 (E.D. Tex. 1956). Sheets v. Yamaha Motors Corp., USA, 657 F.Supp. 319 (1987), Commonwealth v. Robinson, 7 Mass.App.Ct. 470, 388 N.E.2d 705 (1979), MBL (USA) Corp. v. Diekman, 112 Ill.App.3d 229, 445 N.E.2d 418, 67 Ill.Dec. 938 (1983), A. H. Emery Co. v. Marcan Products Corporation, 380 F.2d 11 (1968), Ecologix, Inc. v. Fansteel, Inc., 676 F.Supp. 1374 (1988), Laser Industries, Ltd. v. Eder Instrument Co., 573 F.Supp. 987 (1983), Sandlin v. Johnson, 152 F.2d 8 (8th Cir.1945), Trandes Corp. v. Guy F. Atkinson Co., 996 F.2d 655 (4th Cir.1993), Ferranti Electric, Inc. v. Harwood, 43 Misc.2d 533, 251 N.Y.S.2d 612 (1964), The Boeing Company v. Sierracin Corporation, 108 Wash.2d 38, 738 P.2d 665 (1987). Note that the analysis into factors is mine, and for the purpose of illustration in this paper. It should not be relied on in a court of law.

Table 3. Cases with Factors Grouped by Issue.

Case	P or D	Trade Secret		Misappropriation		
		Maintain Secrecy	Info Valuable	Improper Means	Confidential Relationship	Info Used
Deceived	P	F6p F10d	F24d	F26p		
NoMeasures	D	F10d	F24d	F2p		
Bribed	?	F6p F10d	F6p F16d	F2p		
Mason	P	F6p F15p	F6p F16d		F1d F21p	
Silfen	D	F6p F11d	F6p		F21p	
Lewis	P	F8p			F1d F21p	
College	P	F15		F26p	F1d	
Arco	P	F16d F20d	F10d			
Sheets	D		F19d F27d			F18p
Robinson	D		F10d F19d	F26p	F1d	F18p
MBL	D	F6p F20d	F4p F5d F6p F10d		F1d F4p F5d F13p	
Prentice	D	F6p F24d	F6p	F3d		
Kinnera-Weed	D	F6p	F6p	F25d	F21p	F25d
Emery	P		F6p F10d		F21p	F18p
Laser	P	F6p	F6p F10d F12p		F1d F21p	F18p
Sandlin	D	F6p F16d	F6p F10d		F1d	
Ecologix	D				F1d F21p F23d	
Trandes	P	F6p	F4p F6p F10p F12p	F1d F21p		
Ferranti	D	F20d		F2p S17d		
Boeing	P	F15p	F10d F12p		F1d F21p	F18d

4.1 Explanation in CATO

CATO organised its explanation in terms of issues. The explanation of the example in [42] would be something like the following (adapted from the explanation of *Mason* given in Figure 2.4 of [3]).

Argument for Plaintiff in *Bribed*.

Plaintiff should win a claim of trade secrets misappropriation. Plaintiff's information is a trade secret and defendant acquired plaintiff's information through improper means.

Plaintiff's information is a trade secret.

In *Bribed*, plaintiff adopted security measures [F6p] This shows that plaintiff took efforts to maintain the secrecy of its information.

The fact that plaintiff disclosed its information to outsiders [F10d] does not preclude a conclusion that plaintiff's information is a trade secret. This is especially so where, as in *Bribed*, plaintiff took security measures to protect the information [F6p]. [*Deceived*]

The fact that plaintiff's information could be ascertained by examining or reverse engineering plaintiff's product [F16d] does not preclude a conclusion that plaintiff's information is a trade secret [*Mason*]. Moreover in *Deceived*, the information was available elsewhere, which is not the case in *Bribed*.

Defendant acquired plaintiff’s information through improper means by bribing an employee [F2p].

NoMeasures is not a counterexample with respect to secrecy being maintained, since in that case, unlike *Bribed*, the plaintiff did not take security measures to protect the secrecy of its information.

Organising by issues has several advantages. First it means that differences relating to uncontested issues are not considered: since the use of improper means is not contested, it does not matter which factor is used to establish it. This prunes the *O1a* branch from the tree in Figure 1. Second when downplaying the genuine distinction in *O1c*, it produces the correct rebuttal (P2c) because this is the factor related to the same issue, and ignores *P2’c*. Also it explicitly cites the preference for F6p in past precedents as the reason for citing the precedent and rejecting the counterexample.

The preference for F6p over F16d is justified by reference to *Mason*. Two points should be made here: the misappropriation in *Mason* involved breach of confidence rather than improper means. This does not matter when using issues, but would provide a distinction for an approach without issues as used in [29] and [41]. This is true of the reason as well as the results model: F2p would have to be included in the reason for *Mason*, since otherwise there would have been no breach of confidence. Using issues to organise our factors enables reasoning with portions of precedents [15], which makes substantially more precedents available to make points. On the results model, even with issues, however, *Mason* would be vulnerable to a distinction since it contains F15p whereas *Bribed* does not, which provides an additional factor to outweigh F16d, suggesting that F6p might not be preferred to F16d on its own. If we consider the decision in *Mason*, however, we read:

We note that absolute secrecy is not required ... for the recipe for Lynchburg Lemonade to constitute a trade secret — ”a substantial element of secrecy is all that is necessary to provide trade secret protection.” Drill Parts, 439 So.2d at 49. The defendants also contend that *Mason*’s recipe was not a trade secret because it could be easily duplicated by others. ... We do not think, however, that this evidence in and of itself could prevent such a conclusion. Rather, this evidence should be weighed and considered along with the evidence tending to show the existence of a trade secret. 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. *KFC Corp. v. Marion-Kay Co.*, 620 F. Supp. 1160 (S.D.Ind. 1985); *Sperry Rand Corp. v. Rothlein*, 241 F. Supp. 549 (D.Conn. 1964).

Here there is no reference to the uniqueness of the product (F15p), and so for the reason model we should take it that the decision in favour of *Mason* indicates that F6p is preferred to F16d on its own, without needing the support of F15p, and so *Mason* is available for use in *Bribed* when the reason model is applied at the issue level. Note, however, that if we are using the distribution of

factors across issues in [27] that was shown in Figure 2, this would require us to consider the preference at the level of the *TradeSecret* issue and that InfoValuable is not required for the information to be a trade secret when F6p is present. We prefer instead to include F6p under InfoValuable as well as MaintainSecrecy: sufficient secrecy measures are considered enough to justify the information being deemed valuable. A factor can relate to more than one issue: for example F25d in Figure 2. Putting F6p under InfoValuable to oppose F16d seems to accord with the decision in *Mason* cited above. This enables us to continue to use the logical model of [19] and [27] which requires *Bribed* to show both that the information was valuable and that efforts to maintain secrecy had been taken. Without F6p being included under InfoValuable this model would fail since that there is no factor to contest the claim that F16d would mean that the information lacked value.

4.2 The Centrality of Issues

From the above discussion we see that organising explanations around issues provides focus and enables irrelevant factors and uncontested issues to be ignored, so avoiding swamping the recipient of the explanation with an excess of information. This accords with the widespread popularity of the Issue-Rule-Application-Conclusion (IRAC) methodology of legal analysis. IRAC is widely taught in law schools⁵, although often there are variants which include an additional item or reorder the items, perhaps beginning with the conclusion. IRAC was advocated for the explanation of outcomes from factor based reasoning in [8]. "Issue spotting" has a long history in AI and Law, dating back to the work of Gardner [23] and pursued by Gordon [24] and [25].

The question arises, however, as to what should be counted as an issue. Issues could be very coarse grained such as *was the information a trade secret?*, or relate to the fine grained abstract factors of the CATO factor hierarchy, such as *did the plaintiff take adequate security measures with respect to the defendant?*. In the next section we will offer a tree of issues in the form of an Abstract Dialectical Framework (ADF) [18] as used for the representation of legal knowledge in [1].

5 An ADF of Issues

Table 4 presents the issues used to decide questions of Trade Secret Misappropriation in the form of an ADF. The decomposition of InfoMisappropriated follows [27] as shown in Figure 2 rather than [19]. Traditionally models of reasoning with legal precedents include the *rule* model and the *balance of factors* model [52].

⁵ For example, City University of New York (<https://www.law.cuny.edu/legal-writing/students/irac-crracc/irac-crracc-1/>) and Elizabeth Haub School of Law at Pace University (<https://academicsupport.blogs.pace.edu/2012/10/26/the-case-of-the-missing-a-in-law-school-you-cant-get-an-a-without-an-a/>). Use of IRAC is advocated by the LexisNexis survival guide for law students available at <https://www.lexisnexis.co.uk/students/law/>.

Table 4. ADF for Trade Secret Misappropriation

Node	Children	Acceptance Conditions	Justification
Trade Secret Misappropriation	InfoTradeSecret InfoMisappropriated	ACCEPT IF InfoTradeSecret AND InfoMisappropriated REJECT	Restatement
InfoTradeSecret	InfoValuable MaintainSecrecy	ACCEPT IF InfoValuable AND MaintainSecrecy REJECT	Restatement
Info Misappropriated	WrongDoing InfoUsed	ACCEPT IF WrongDoing AND InfoUsed REJECT	Restatement
InfoValuable	F6p F8p F11d F15p InfoObtainable	REJECT IF F11d ACCEPT IF F8p ACCEPT IF F6p ACCEPT IF F15p REJECT IF InfoObtainable ACCEPT	Silfen Lewis Mason College Restatement
MaintainSecrecy	F27d F6p F19d MeasuresOutsiders	REJECT IF F27d REJECT IF F19d ACCEPT IF F6p REJECT IF NOT MeasuresOutsiders ACCEPT	Sheets Robinson Emery Restatement
WrongDoing	F3d OwnEfforts ImproperMeans ConfidRelation	REJECT IF F3d REJECT IF OwnEfforts ACCEPT IF ImproperMeans ACCEPT IF ConfidRelation REJECT	Prentice Restatement Restatement Restatement
ImproperMeans	InfoMisuse IllegalAct	ACCEPT IF InfoMisuse ACCEPT IF IllegalAct REJECT	Restatement Restatement
ConfidRelation	NoticeConfid ExplicitAgreement	ACCEPT IF NoticeConfid ACCEPT IF ExplicitAgreement REJECT	Restatement Restatement
InfoUsed	InfoMisuse OwnEfforts F8p F18p	ACCEPT IF InfoMisuse ACCEPT IF F8p ACCEPT IF F18p REJECT IF OwnEfforts ACCEPT	Restatement Lewis Emery Restatement
InfoMisuse	F7p F14p	ACCEPT IF F7p OR F14p	Restatement
IllegalAct	F2p F22p F26p	ACCEPT if F2p OR F22p OR F26p ACCEPT	Restatement
NoticeConfid	F1d F21p F23d	REJECT if F23d ACCEPT IF F21p REJECT IF F1d ACCEPT	Ecologix Laser Sandlin
ExplicitAgreement	F4p F5d F13p	REJECT IF F5d ACCEPT IF F4p OR F13p REJECT	MBL Trandes
InfoObtainable	F15p F16d F20d F24d	REJECT IF F15p ACCEPT IF F24d OR F20d REJECT	College Ferranti
MeasuresOutsiders	F10d F12p	ACCEPT IF F12p REJECT IF F10d ACCEPT	Trandes Arco
OwnEfforts	F17d F25d	ACCEPT IF F17d OR F25d REJECT	Kinnear-Weed

Although CATO style reasoning reflects the balance of factors model, the ADF in Table 4 uses the rule model in all the acceptance conditions, assuming that there are sufficient precedents to justify the priorities in every node. In Table 4 the fourth column shows the justification for the acceptance condition: either the logical model implied by the *Restatement* or a precedent which expressed a particular preference⁶.

However, following the examples of CABARET [51] and Issue Based Prediction (IBP) [19], at some point in each branch the rule model might not be appropriate and so the balance of factors model would be required. This will be particularly so in the early stages of the development of a body of case law,

⁶ Because the use of issues mean that at most five factors need to be considered for any given node, it is feasible to envisage enough precedents to resolve the node. This would not be so as the whole case level, where there are 2^{26} possible models.

when there will not yet be sufficient precedents to determine all the required preferences. For example, F18p (Unique Product) may be a strong indication that the information was used, but might not be as decisive as suggested in the ADF in Table 4. If we use a balance of factors model we can choose between balancing the factors representing the children of the node in question or the factors which result from unfolding the nodes below the node concerned. Thus InfoTradeSecret could be resolved by regarding InfoValuable and InfoTradeSecret as a conjunction so that both must be true (rule model); or allowing some trade off between them (coarse grained balance of factors), or balancing the relevant base level factors {F6p, F8p, F10d, F12p, F11d, F15p, F16d, F19d, F20d, F24d, F27d} (fine grained balance of factors). Also perhaps a mixed granularity could be used, grouping some of the base level factors into abstract factors, e.g. F10d and F12p could be considered together as MeasuresOutsiders. It would be an interesting exercise to see what effect these different granularities might have, but for this paper we will assume a mature domain allowing us to use a pure rule model.

5.1 Using Issues for Explanation

Now we have the ADF, we can identify the issues in particular cases. The issues in a case will be the lowest nodes spanning both a pro-plaintiff factor and a pro-defendant factor. We illustrate this principle by applying it to the cases in Table 2:

In *Deceived* MaintainSecrecy is an issue because Security Measures (F6p) favours the plaintiff whereas MeasuresOutsiders favours the defendant (F10d). *Deceived* was found for the plaintiff, and so MeasuresOutsiders \prec F6p. We could use *Emery* as a precedent to justify this preference. InfoValuable is also an issue because it contains both F6p and F24d. Again since the plaintiff won we can infer that InfoAvailableElsewhere \prec F6p, a preference justified by *Mason*.

In IRAC terms:

An issue is whether that there were no efforts to maintain secrecy with respect to outsiders means that secrecy was not maintained. The rule is that Not MeasuresOutsiders \prec F6p (*Emery*). Application is that Not MeasuresOutsiders applies because F10d is present, but F6p is also present. Therefore secrecy was maintained. A second issue is whether that the information was obtainable elsewhere means that the information was not valuable. The rule is InfoAvailableElsewhere \prec F6p (*Mason*). The rule applies because F24d establishes InfoAvailableElsewhere and F6 is present. Therefore, the information was valuable.

In *NoMeasures* only the root node, TradeSecretMisappropriation, spans contested factors. Here the defendant can establish that the information is not a trade secret because it wins both branches for that issue and so the defendant is found for.

In IRAC terms:

The issue is whether a Trade Secret was misappropriated, when the information was misappropriated but not a trade secret. The rule is if not InfoTradeSecret then not TradeSecretMisappropriation (*Restatement of Torts*). The rule applies because the information was obtainable elsewhere because F24d was present (*Ferranti*) and the efforts to maintain secrecy with respect to outsiders were inadequate because F10d is present (*Arco*). Therefore there was no TradeSecretMisappropriation.

In *Bribed* we have a situation similar to *Deceived* except that InfoAvailableElsewhere is established by F16d rather than F24d. The IRAC explanation is thus similar.

An issue is whether in there were inadequate efforts to maintain secrecy with respect to outsiders means that secrecy was not maintained. The rule is that MeasuresOutsiders \prec F6p (cf *Deceived*). Application is that MeasuresOutsiders applies because F10d is present, but F6p is also present. Therefore secrecy was maintained. A second issue is whether that the information was obtainable elsewhere means that the information was not valuable. The rule is InfoAvailableElsewhere \prec F6p (*Mason*). The rule applies because F16d establishes InfoAvailableElsewhere and F6 is present. Therefore, the information was valuable.

In *Mason* there are two issues. The first, InfoValuable, is the what sets our precedent for InfoAvailableElsewhere \prec F6p. The second relates to whether there was notice of confidentiality with both F1d and F21p present. The answer is that there was because F1d \prec F21p, established in *Laser*.

In IRAC terms:

An Issue is whether that the information was obtainable elsewhere means that the information was not valuable. The rule used by this court is InfoAvailableElsewhere \prec F6p. The rule applies because F26d establishes InfoAvailableElsewhere and F6 is present. Therefore, the information was valuable. A second issue is whether there was notice of confidentiality where information was disclosed in negotiations and the defendant knew the information to be confidential. The rule is F1d \prec F21p (*Laser*). The rule applies because F1d and F21p are present. Therefore there was notice of confidentiality.

Whereas the explanations from [3] and [42] begin at the top level and work down to the decisive facts, the IRAC explanations begin with the decisive facts. The IRAC explanations are thus very focused and do not aspire to give an exhaustive account, dotting every 'i' and crossing every 't', but instead home in on what mattered in the particular case. As such they assume that the person to whom the explanation will have some knowledge of the domain. A person familiar with the logical model in Figure 2 or, better yet, the nodes of the ADF in Table 4 will have no difficulty in seeing why these points matter and how they decide the case. This suppression of shared background to highlight the decisive considerations was one of the original motivations for argument based

explanation [13]. For those who need or want a fuller explanation, a dialogue seeking summary information in the manner of [12] can be initiated. Thus asking *SO?* will display the parent node of the issue and asking *WHY?* will display a child node.

Thus for *Bribed* asking *SO?* for the first issue will produce *Secrecy was Maintained*. Asking *SO?* again will produce *The Information was a Trade Secret*. A third *SO?* will produce *The Trade Secret was Misappropriated*. Now a series of *WHY?*s will produce *The Information Was Misappropriated*, *There Was Wronggoing*, *There was an Illegal Act* and finally *The Information Was Obtained by Deception*. Of course the users may stop the flow of information when they have enough to see the correctness of the solution.

5.2 The Boeing Company v. Sierracin Corporation, 108 Wash.2d 38, 738 P.2d 665 (1987)

We will now give a full example, using *Boeing*⁷. Boeing is a very well known aircraft manufacturer. They do not, however, make all the parts themselves, but sub-contract certain components. One such component was the 7/7/7 cockpit windows. Five of these windows lie on each side of the aircraft's nose, each performing multiple critical functions such as defogging, withstanding cabin pressurisation, and providing a clear range of vision. The three major suppliers of aircraft windows in the United States are PPG Industries, Inc.; Swedlow, Inc.; and Sierracin. Sierracin had supplied Boeing with other products for many years. Boeing's drawings for the 7/7/7 cockpit windows are unique, detailed blueprints containing approximately 500 critical tolerances, dimensions, specifications and material requirements. Boeing helped Sierracin enter the 7/7/7 window market by providing Sierracin with FAA authorized drawings, technical assistance and tooling, and by awarding it contracts for some of Boeing's 7/7/7 window needs in 1982 and 1983.

In 1984 after alleged breaches of contract because of late deliveries of windows, Boeing chose not to renew contracts with Sierracin, and instead signed a 5-year 100 percent requirements contract with PPG Industries, Inc. Sierracin decided, however, to continue manufacturing windows for sale on its own in the 7/7/7 "after market" (i.e., spare parts market). As a supplier, Sierracin received Boeing's requests for quotations, which provided that all orders were subject to its confidential terms and conditions. Boeing alleged that Sierracin misappropriated its trade secrets concerning the design of aeroplane windows.

Sierracin signed over 270 contracts with Boeing, each containing the following language: "Confidential Disclosure. Seller shall keep confidential . . . all proprietary information".

The defence was that the information was not a trade secret because it had been disclosed to outsiders (F10d) and had not been misappropriated because it had been disclosed to Sierracin in negotiations (F1d). This was countered by the

⁷ The description of the case is based on the opinion of Justice Dore, available at <https://casetext.com/case/boeing-company-v-sierracin-corporation>.

explicit confidentiality agreement that Boeing required suppliers to sign (F12p), and that Sierracin were well aware of the confidential nature of the information (F21p). The court found for Boeing.

The explanation based on the above proposal would look like this:

The Boeing Company v. Sierracin Corporation, 108 Wash.2d 38, 738 P.2d 665 (1987).

The decision is for the plaintiff. There are two issues:

1. Whether adequate measures with respect to outsiders were taken when the information was disclosed to outsiders, but these disclosures were restricted. The rule is Secrets-Disclosed-Outsiders \prec Outsider-Disclosures-Restricted (Trandes Corp. v. Guy F. Atkinson Co., 996 F.2d 655 (4th Cir.1993)). The rule applies because F10d and F12p are present. Therefore, adequate measures with respect to outsiders were taken.
2. Whether there was notice of confidentiality when the information was disclosed in negotiations, but the defendant knew that the information was confidential. The rule is Disclosure-In-Negotiations \prec Knew-Info-Confidential (Laser Industries,Ltd. v. Eder Instrument Co., 573 F.Supp. 987 (1983)). The rule applies because F1d and F21p are present. Therefore, there was notice of confidentiality,

The user may now interrogate the system further to help in understanding why these issues matter.

Issue 1: So?

Reply 1: Secrecy was Maintained (Restatement of Torts section 757, comment(b), bullet 3).

Reply 1: So?

Reply 2: The information was a Trade Secret (Restatement of Torts section 757, comment(b)).

Reply 2: Why?

Reply 3: The information was valuable. (Restatement of Torts section 757, comment(b), bullet 3).

Reply 3: Why?

The issue was unopposed. Further the product was unique (F15p).

Satisfied as to Issue 1, the user now turns to issue 2.

Issue 2: So?

Reply 4: There was a Confidential Relationship (Restatement of Torts, section 757(b)).

Reply 4: So?

Reply 5: The Information Was Misappropriated (Restatement of Torts, Section 757, General Principle).

OK

Since the user knows that the plaintiff should win if the information was a trade secret and misappropriated, the dialogue is terminated here.

6 Discussion

In this section we will consider three points: the quality of the explanations; the implications for accounts of precedential constraint; and how to accommodate dimensional facts.

6.1 Quality of Explanation

In his illuminating survey on explanation [39], Miller identifies four features of good explanations. These are:

- Explanations are *contrastive*. As well as explaining why a particular classification is appropriate, a good explanation will also say why other classifications are not, often using counterfactuals and hypotheticals.
- Explanations are *selective*. Rarely is a logically complete explanation provided, but rather only the most salient points are presented unless more detail is required by the recipient of the explanation. The assumption is that there will be a considerable degree of shared background knowledge, and so the explanation need only point to some fact or rule as yet unknown to the recipient.
- Explanations are rarely in terms of *probabilities*. Using statistical generalisations to explain why events occur is unsatisfying since they do not explain the generalisation itself. Moreover, the explanation typically applies to a single case, and so would require some explanation of why that particular case is typical.
- Explanations are *social*. Explanations involve a transfer of knowledge, between particular people in a particular situation and so are relative to the explainer’s beliefs about the explainee’s beliefs.

The explanations produced by using issues as described above have these features. The explanation is contrastive because it begins with the issue which will include a factor for the other side, and so suggests why the decision would have gone otherwise had the preferred factor for the winning side not been present. They are selective because they begin by stating the particular reason for the decision, and offers any further explication of the background knowledge explaining exactly why this matters only on request. Like most explanations in law, no probabilities are used: specific precedents are used to justify the rules rather than some degree of support as in, for example, association rule mining [53]. Finally the explanation is social in that it adopts a form of explanation (IRAC) that is widely used in the legal community which transfers knowledge in a way tailored to the situation and the user. The particular strength of the proposed method is perhaps its *selectivity*, which contrasts with the dialogue proposed in [42] which included all available arguments and objections, and the explanations in [3] and [1] which again covered every aspect without regard as to what the user might already know.

6.2 Implications for Precedential Constraint

The increased effectiveness of formal characterisations of precedential constraint when applied at the issue level rather than at the whole case level was discussed in [11]. A problem with the results model was that, given, the large, number of factors, there are so many possible case descriptions (2^{26}) that it is all too easy to avoid the constraint of a precedent by pointing to a distinguishing factor. The use of the reason model alleviates this problem to some extent, but far from completely, as illustrated in [11]. However, the number of possible case descriptions at the *issue* level is greatly reduced. Examination of Table 4 shows that no node has more than five children and so no node would require more than 32 cases to be fully resolved. Many of the nodes would require fewer: MeasuresOutsiders, measures for example has only two children, so that there are only four distinct case descriptions. Actually only three of these are possible, since F12p (restrictions placed on outsider disclosures) cannot be present without F10d (disclosures to outsiders). Also some precedents can be used for several nodes: the cases of MeasuresOutsiders required in MaintainSecrecy can also be used to resolve MeasuresOutsiders itself. This being so, fewer than 150 precedents would be required to resolve the tree completely, even on the results model. This number could be substantially reduced by applying the reason model to the issues. Since IBP [19] used over 180 cases, it would seem that this might be a feasible way of addressing the problem. Moreover, various Machine Learning approaches use even bigger datasets: over 15,000 were available to [17]. If we have a dataset already available in a form in which factors can be straightforwardly ascribed - as might be expected when decisions are made on the basis of an application form, which is common in fields such as welfare benefits - or we have a machine learning program which ascribes factors as in [17], we can apply this approach if we can associate these factors with issues. This should enable us to learn the acceptance conditions for the nodes using a variety of ML techniques, including such traditional techniques as rule induction [40].

The above suggests that whether we are using precedents for explanation or for learning, they are best considered in terms of issues rather than as whole cases.

6.3 Factors With Magnitude

The above presents a rather sanguine view of the possibility of predicting legal decisions, suggesting that a couple of hundred cases should enable the identification of a set of rules that would give complete accuracy. Leaving aside the many problems with any approach predicting legal decisions on the basis of a set of past cases, including that the law is constantly evolving so that predictions based on old data become unreliable [37] and that any collection of precedents is likely to contain decisions that were biased or incorrect, there is another serious problem. The above has assumed that, as in [3], factors are either present or absent, and can be ascribed to cases unequivocally. In practice, however, this is not so. Consider the Restatement of Torts:

Some factors to be considered in determining whether given information is one's trade secret are: (1) the **extent** to which the information is known outside of his business; (2) the **extent** to which it is known by employees and others involved in his business; (3) the **extent** of measures taken by him to guard the secrecy of the information; (4) the **value** of the information to him and to his competitors; (5) the **amount** of effort or money expended by him in developing the information; (6) the **ease or difficulty** with which the information could be properly acquired or duplicated by others. **Emphasis mine.**

From this it is clear that many of the aspects are not simply present or absent, but are present or absent to some degree and so require some judgement as to whether they were present to a degree *sufficient* to permit the ascription of the factor. It should be remembered that the factors used in CATO derive from the dimensions proposed in HYPO ([47] and [4]). Dimensions were aspects of a case which could, if applicable, take a range of values which would increasingly favour a particular party. The relationship between dimensions and factors is discussed in [48] and [11]. In fact, in HYPO, ten of the thirteen dimensions could take only two values and so either were inapplicable or favoured one of the parties. For example, the bribery dimension either favoured the plaintiff if bribery had taken place, or was inapplicable. It therefore maps straightforwardly to a single factor, F2p. Three of the dimensions did, however, span a range of values. Competitive Advantage was either inapplicable, neutral, or favoured the plaintiff, mapping to F8p. Disclosures to Outsiders was inapplicable if there had been no disclosures, or favoured the defendant if there had been sufficient disclosures (F10d) and neutral otherwise. Note, however, that extreme pro-defendant values on this dimension gave rise to the more powerful factor F27d when the information was considered in the public domain. This is important because a plaintiff factor might be preferred to F10d, but not F27d. This is true of F6p in Table 4. The most interesting dimension is security measures. At one end this favours the defendant and so maps into F19d, whereas at the other it maps into F6p and favours the plaintiff. It is thus always applicable. Many cases, however, contain neither F19d nor F6p, suggesting that the middle of the range is neutral so that no factor is applicable. Note that where a dimension can favour both sides, two distinct factors are used: this is because in general the absence of factor is not a reason to decide for the other side. This was explained by Rissland and Ashley in [48]:

[...] the Security-Measures dimension was broken into two factors: Security-Measures [F6p], favoring the plaintiff, and No-Security-Measures [F19d], favoring the defendant. This was done because judges explicitly said that the fact that plaintiff had taken no security measures was a positive strength for the opponent. By contrast, Ashley and Alevén did not create a “No-Secrets-Disclosed-Outsiders” factor because they found no cases where judges had said that the absence of any disclosures to outsiders was a positive strength for the plaintiff. ([48], p 69).

Most of the CATO factors derive from two valued dimensions and so can be considered either present or absent, and so do not require special consideration. The three factors deriving from dimensions with ranges of value, however, do need something more. Horty in [30] and [31] introduced the notion of *factors with magnitude*, and discussed how these could be accommodated in a theory of precedential constraint. Rigoni addressed this problem in [46], and Horty modified his approach in [32]. A comparison of their approaches is given in [42]. Although this was taken in [42] to imply that precedential constraint should be expressed in terms of dimensions rather than factors, it was argued in [11] that this is perhaps not the best approach. We will explain how we can accommodate dimensions in an account of precedential constraint based on factors.

Many have seen reasoning with legal cases as a two stage process: first factors are assigned on the basis of the facts in a cases, and then these factors are considered in the the light of precedent cases to see whether the decision is constrained. This two stage approach is described in [44]

Once the facts of a case have been established - and this is rarely straightforward since the move from evidence to facts is often itself the subject of debate - legal reasoning can be seen, following Ross [49] and Lindhal and Odelstad [36], as a two stage process, first from the established facts to intermediate predicates, and then from these intermediate predicates to legal consequences. CATO has been explicitly identified with the second of these steps (e.g. [19]). ([44], p 22).

This approach has been used not only in [44], but also in [5] and further advocated in [17] and [16].

If we adopt this two stage model, we can see factors with magnitude as factors deriving from a dimension with more than two values. It may, like Security Measures, favour either side, with a neutral area in which no factor is applicable, or, like disclosures, give rise to two factors favouring the same side with different strengths, as with F10p and F27p. Now we must determine which factor, if any, should be ascribed in a particular case given that it lies at a certain point on the range (termed by Horty a *dimensional fact*). This can be determined by precedents using either the results model or the reason model [31]. Rigoni proposed that precedents should be regarded as identifying *switching points* on the dimensions, the points at which factors come to be, and cease to be, applicable. The fact that whether a factor is applicable or not is itself be debatable, complicates the first stage of the process. However, once the set of applicable factors has been identified, the second stage can proceed as described above.

Another, perhaps more serious, problem is that these factors may not be independent. For example, cases arise which require balancing the interests of the state in enabling the enforcement of laws with the privacy interests of its citizens [14]. In such cases it may seem necessary to trade off one factor against another, so that the more serious the suspected crime the greater the intrusion on privacy that is justified. Such balancing of interests has been discussed in [35] and [26]. We would, however, in line with the two stage approach, follow [10]




Fig. 3. Trade off between absence and income

and see the question not as a balance between factors, but as a question of the ascription of factors on the basis of dimensional facts, so that factors themselves can continue to be seen as independent.

The example in [31] concerns change of fiscal domicile. Among other things to be considered are the length of stay abroad and the percentage of income earned abroad. The longer the absence and the greater the amount, the more change is favoured. We may now have a decision where an absence of 36 months and earnings of 60% favoured change, while an absence of 48 months but only 20% earnings favoured no change. This suggests that absence and income are not independent, but trade off against each other. Suppose we have a third case also with 20% earnings, but an absence of 60 months, further indicating the existence of a trade off.

The suggestion in [10] is to introduce a factor ascribed to the case on the basis of the dimensional *facts* recording absence and income. In this case a suitable factor would be *IncomeSufficientGivenAbsence* and would favour change. Each precedent for change will block off an area where the factor definitely applies, and each precedent for the defendant will block off an area where the factor definitely does not apply. This is shown in Figure 3. We can now fit a line to the points and suggest that the factor applies to points north east of the line and does not apply to points south west of the line. A possible line ($y = 120 - 10x$: designed to just include both precedents) is shown in Figure 3. Of course, other lines are possible, and the function need not be linear, and so any unconstrained point may be the subject of debate as to whether or not the precedent applies.

The composite factor will appear as a node in the ADF with the dimensional facts as its children and the equation will define the acceptance condition. This node can be treated as an issue for explanation purposes.

7 Concluding Remarks

In this paper we have discussed several reasons why cases are better seen as bundles of issues than as bundles of factors. The issues will be the bundles of

factors. Thinking in terms of issues improves explanations by enabling them to focus on what was disputed and what is significant in the particular case under consideration, and to be expressed in the IRAC form widely taught in law schools. Using issues also greatly enhances factor based precedential constraint by eliminating irrelevant distinctions. The importance of issues for prediction was indicated by the central role given to issues in systems designed to predict legal decisions based on factors such as [19] ad [28], and is discussed in detail in [11].

Developing good factor based explanations is of current importance because of the increased use of machine learning approaches for predicting legal decisions. Explanations are essential because of the right to explanation and to encourage acceptance of these predictions. But they also have relevance to the responsible use of such systems: machine learning approaches are vulnerable to changes in the law and social attitudes, and the bias that may exist in the past decisions. Good explanations will help to detect decisions based on reasons which are no longer applicable, and decisions based on reasons that exhibit bias. Explanation can therefore apply a corrective influence essential to the responsible use of Machine Learning in law.

References

1. Al-Abdulkarim, L., Atkinson, K., Bench-Capon, T.: A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks. *AI and Law* **24**(1), 1–49 (2016)
2. Aletras, N., Tsarapatsanis, D., Preotiuc-Pietro, D., Lampos, V.: Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science* **2**: DOI: [10.7717/peerj-cs.93](https://doi.org/10.7717/peerj-cs.93) (2016)
3. Aleven, V.: Teaching case-based argumentation through a model and examples. Ph.D. thesis, University of Pittsburgh (1997)
4. Ashley, K.D.: Modeling legal arguments: Reasoning with cases and hypotheticals. MIT press, Cambridge, Mass. (1990)
5. Ashley, K.D., Brüninghaus, S.: Automatically classifying case texts and predicting outcomes. *AI and Law* **17**(2), 125–165 (2009)
6. Atkinson, K., Bench-Capon, T., Bollegala, D.: Explanation in AI and Law: Past, present and future. *Artificial Intelligence* p. 103387 (2020)
7. Bench-Capon, T.: HYPO’s legacy: Introduction to the virtual special issue. *AI and Law* **25**(2), 1–46 (2017)
8. Bench-Capon, T.: Explaining legal decisions using IRAC. In: *Computational Models of Natural Argument 2020*. vol. 2669, pp. 74–83. CEUR Workshop Proceedings 2669 (2020)
9. Bench-Capon, T.: The need for good old fashioned AI and Law. In: Hötendorfer, W., Tschohl, C., Kummer, F. (eds.) *International trends in legal informatics: a Festschrift for Erich Schweighofer*, pp. 23–36. Weblaw, Bern (2020)
10. Bench-Capon, T., Atkinson, K.: Dimensions and values for legal CBR. In: *Proceedings of JURIX 2017*. pp. 27–32. IOS Press (2017)
11. Bench-Capon, T., Atkinson, K.: Precedential constraint: The role of issues. In: *Proceedings of the 18th International Conference on AI and Law*. p. In Press. ACM (2021)

12. Bench-Capon, T., Coenen, F., Orton, P.: Argument-based explanation of the British Nationality Act as a logic program. *Information and Communications Technology Law* **2**(1), 53–66 (1993)
13. Bench-Capon, T., Lowes, D., McEnery, A.: Argument-based explanation of logic programs. *Knowledge-Based Systems* **4**(3), 177–183 (1991)
14. Bench-Capon, T., Prakken, H., Visser, W.: Argument schemes for two-phase democratic deliberation. In: *Proceedings of the 13th International Conference on AI and Law*. pp. 21–30 (2011)
15. Branting, L.K.: Reasoning with portions of precedents. In: *Proceedings of the 3rd International Conference on AI and Law*. pp. 145–154 (1991)
16. Branting, L.K.: Explanation in hybrid, two-stage models of legal prediction. In: *The 3rd XAILA Workshop at JURIX 2020* (2020)
17. Branting, L.K., Pfeifer, C., Brown, B., Ferro, L., Aberdeen, J., Weiss, B., Pfaff, M., Liao, B.: Scalable and explainable legal prediction. *AI and Law* (Available Online 2020)
18. Brewka, G., Woltran, S.: Abstract Dialectical Frameworks. In: *Twelfth International Conference on the Principles of Knowledge Representation and Reasoning* (2010)
19. Brüninghaus, S., Ashley, K.D.: Predicting outcomes of case based legal arguments. In: *Proceedings of the 9th International Conference on AI and Law*. pp. 233–242. ACM (2003)
20. Chalkidis, I., Androutsopoulos, I., Aletras, N.: Neural legal judgment prediction in English. arXiv preprint arXiv:1906.02059 (2019)
21. Chorley, A., Bench-Capon, T.: An empirical investigation of reasoning with legal cases through theory construction and application. *AI and Law* **13**(3-4), 323–371 (2005)
22. Doshi-Velez, F., Kortz, M., Budish, R., Bavitz, C., Gershman, S., O’Brien, D., Schieber, S., Waldo, J., Weinberger, D., Wood, A.: Accountability of AI under the law: The role of explanation. arXiv preprint arXiv:1711.01134 (2017)
23. Gardner, A.V.D.L.: *An artificial intelligence approach to legal reasoning*. MIT Press (1987)
24. Gordon, T.F.: Issue spotting in a system for searching interpretation spaces. In: *Proceedings of the 2nd international conference on AI and Law*. pp. 157–164 (1989)
25. Gordon, T.F.: The pleadings game. *AI and Law* **2**(4), 239–292 (1994)
26. Gordon, T.F., Walton, D.: Formalizing balancing arguments. In: *Proceedings of COMMA 2016*. pp. 327–338 (2016)
27. Grabmair, M.: *Modeling Purposive Legal Argumentation and Case Outcome Prediction using Argument Schemes in the Value Judgment Formalism*. Ph.D. thesis, University of Pittsburgh (2016)
28. Grabmair, M.: Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs. In: *Proceedings of the 16th International Conference on AI and Law*. pp. 89–98 (2017)
29. Horty, J.F.: Reasons and precedent. In: *Proceedings of the 13th International Conference on AI and Law*. pp. 41–50 (2011)
30. Horty, J.F.: Reasoning with dimensions and magnitudes. In: *Proceedings of the 16th the International Conference on AI and Law*. pp. 109–118 (2017)
31. Horty, J.F.: Reasoning with dimensions and magnitudes. *AI and Law* **27**(3), 309–345 (2019)
32. Horty, J.F.: Modifying the reason model. *AI and Law* (Available On Line 2021)
33. Horty, J.F., Bench-Capon, T.: A factor-based definition of precedential constraint. *AI and Law* **20**(2), 181–214 (2012)

34. Kaur, A., Bozic, B.: Convolutional neural network-based automatic prediction of judgments of the european court of human rights. In: 27th AIAI Irish Conference on AI and Cognitive Science. pp. 458–469. CEUR 2563 (2019)
35. Lauritsen, M.: On balance. *AI and Law* **23**(1), 23–42 (2015)
36. Lindahl, L., Odelstad, J.: Open and closed intermediaries in normative systems. In: Proceedings of JURIX 2006. pp. 91–99. IOS Press (2006)
37. Medvedeva, M., Vols, M., Wieling, M.: Using machine learning to predict decisions of the European Court of Human Rights. *AI and Law* pp. 1–30 (2019)
38. Medvedeva, M., Xu, X., Vols, M., Wieling, M.: JURI says: An automatic judgement prediction system for the European Court of Human Rights. In: Proceedings of JURIX 2020. pp. 277–280 (2020)
39. Miller, T.: Explanation in Artificial Intelligence: Insights from the social sciences. *Artificial Intelligence* **267**, 1–38 (2019)
40. Možina, M., Žabkar, J., Bench-Capon, T., Bratko, I.: Argument based machine learning applied to law. *AI and Law* **13**(1), 53–73 (2005)
41. Prakken, H.: A formal analysis of some factor- and precedent-based accounts of precedential constraint. *AI and Law* (Available On-Line 2021)
42. Prakken, H., Rosa, R.: A top-level model of case-based argumentation for explanation: formalisation and experiments. *Argument and Computation* (Available On-Line 2021)
43. Prakken, H., Sartor, G.: Modelling reasoning with precedents in a formal dialogue game. *AI and Law* **6**(2-4), 231–287 (1998)
44. Prakken, H., Wyner, A., Bench-Capon, T., Atkinson, K.: A formalization of argumentation schemes for legal case-based reasoning in ASPIC+. *Journal of Logic and Computation* **25**(5), 1141–1166 (2015)
45. Rigoni, A.: An improved factor based approach to precedential constraint. *AI and Law* **23**(2), 133–160 (2015)
46. Rigoni, A.: Representing dimensions within the reason model of precedent. *AI and Law* **26**(1), 1–22 (2018)
47. Rissland, E.L., Ashley, K.D.: A case-based system for Trade Secrets law. In: Proceedings of the 1st International Conference on AI and Law. pp. 60–66. ACM (1987)
48. Rissland, E.L., Ashley, K.D.: A note on dimensions and factors. *AI and Law* **10**(1-3), 65–77 (2002)
49. Ross, A.: Tû-tû. *Harvard Law Review* **70**, 812–825 (1957)
50. Schauer, F.: Precedent. *Stanford Law Review* **39**, 571–605 (1987)
51. Skalak, D.B., Rissland, E.L.: Arguments and cases: An inevitable intertwining. *AI and Law* **1**(1), 3–44 (1992)
52. Varsava, N.: How to realize the value of stare decisis: Options for following precedent. *Yale Journal of Law and the Humanities* **30**, 62–120 (2018)
53. Wardeh, M., Bench-Capon, T., Coenen, F.: Padua: a protocol for argumentation dialogue using association rules. *AI and Law* **17**(3), 183–215 (2009)
54. Wyner, A., Bench-Capon, T.: Argument schemes for legal case-based reasoning. In: Proceedings of JURIX 2007. pp. 139–149 (2007)
55. Zheng, H., Grossi, D., Verheij, B.: Case-based reasoning with precedent models: Preliminary report. In: Proceedings of COMMA 2020. pp. 443–450 (2020)