Statement Types in Legal Argument

Latifa AL-ABDULKARIM [a,1], Katie ATKINSON [a] Trevor BENCH-CAPON [a]
[a] *Department of Computer Science, University of Liverpool, UK*

Abstract. In this paper we present an overview of the process of argumentation with legal cases, from evidence to verdict. We identify the various different types of statement involved in the various stages, and describe how the various types relate to one another. In particular we show how we can derive the legally accepted facts which form the basis for consideration of the law governing the cases from facts about the world. We also explain how we can determine which particular facts are relevant. In so doing we bring together several important pieces of AI and Law research and clarify their relationships.

Keywords. legal reasoning, cases, factors, dimensions, evidence

1. Introduction

Modelling argumentation based on legal cases has been a central topic of AI and Law since its very beginnings and several well known systems have contributed to our understanding of this sort of argumentation. The process begins with evidence and ends with the decision of the court, and passes through several stages in between. Although there have been previous attempts to identify different levels or stages in the reasoning (e.g. [7]), the majority of approaches in this area deal with only part of the process. In this paper we will attempt to give the whole picture, from evidence to verdict, focusing in particular on the transitions between stages, and the different types of statement used at the various stages. In particular we need to account for:

- The different types of statement used in legal reasoning. Conflating these different types can make deciding their acceptability problematic.
- The transition between reasoning about facts to reasoning about law. The importance of this transition was noted in [15], but the vast majority of work has been on one side or the other, and how this barrier is crossed remains an open question;
- How and where uncertainty is dealt with. Although many legal concepts are characterised by *open texture* [31] and so have a *penumbra of uncertainty* [21], the verdict itself can permit no qualification. How and where the doubts are resolved remains an open question. Most approaches (e.g. [26]) represent legal concepts only as Booleans. For example factors are considered present or absent and values are promoted or demoted, whereas we believe that differing degrees of presence and absence (e.g. [11], [18], [1]) need to be recognised.

[1] Corresponding Author: T. Bench-Capon. E-mail: tbc@csc.liv.ac.uk

2. Layers of Legal Reasoning

The current established view is that first, facts must be determined on the basis of evidence. Next, these facts must be used to ascribe legally significant predicates to the case on the basis of which the outcome can be established ([24]). In this section we will further articulate this progression, referring to a number of landmark AI and Law systems. We will proceed top-down, starting with the *verdict*. Like most legal CBR, we have in mind US civil cases, based on common law. In criminal law there is more focus on what the evidence establishes, and in European civil law precedents are given less emphasis.

2.1. Verdict

At the top level we have the *verdict* of the court. The verdict can be regarded as a performative utterance in the sense of [9]. It is not true or false: a pronouncement by the appropriate person that he or she finds for the plaintiff simply is how a case is decided for the plaintiff: the pronouncement of the verdict makes it so. As such, these nodes have assertability (sometimes called *felicity*) rather than acceptance conditions. Of course, coming from an AI system, the verdict cannot be a performative, and rather corresponds to a recommendation to the court or a prediction of what the court will do. This latter, which follows the Issue Based Prediction system (IBP) [17], is perhaps the best way to categorise the aims of such systems. Not all AI and Law systems have sought to make recommendations or predictions: HYPO [5], the system from which so much work in AI and Law on arguing with cases stems, was designed to find, but not to evaluate, arguments and its child CATO [3] was specifically designed to instruct law students in a particular kind of argument, namely distinguishing cases, and emphasising or downplaying these distinctions.

The verdict is binary: the verdict must be "yea or nay": the judge cannot refuse to answer the question, nor express uncertainty or doubt as to the verdict. Leave to appeal may be given, but the verdict remains in force until it is overturned or quashed on appeal. The court is required to justify the verdict. Justification is in terms of issues. The relation between issues and verdict, is, as is agreed by all of [29], [17] and [26], deductive: the verdict is expected to follow logically once the issues have been resolved. Perhaps the clearest treatment of the relation between verdict and issues is found in the so-called *logical models of legislation system*s, most sharply presented in [28]. In those systems there is only the logical model: everything below what is explicit in legislation is left to the user for resolution.

The logical model may come from several sources: from statute, as in [28] and [29], from a commentary or other summary of the common law as in [17], or from a precedent case as in [26]. It might at first sight appear that the logical model takes us little further forward: the model in [17] begins by telling us only that the court can find that a trade secret was misappropriated if and only if there was a trade secret and it was misappropriated. But when we go down a level we can find how to establish the existence of a trade secret (it must be of value, and efforts must have been taken to maintain secrecy), and how we can establish misappropriation (either by showing that the information was obtained improperly, or that a confidential relationship was breached by the use of the information). Thus the logical model tells us precisely what arguments can be presented, what *issues* are relevant, how a case should be presented in terms of

these issues, and the consequences of resolving these issues in particular ways. Once the issues have been resolved, the verdict can be justified in purely deductive terms: it is in the resolution of the issues that the distinctively legal argumentation is encountered.

2.2. Issues

As we saw in the last section the verdict is logically entailed by the resolution of the issues forming the logical model. The leaf issues are where we pass from an area of discussion which gives degrees of strength to legal reasons to an area where various propositions are held to be definitively true or false: from CATO style factor based reasoning to standard deduction. It is here that we must decide whether particular issues favour the plaintiff or the defendant, and it is not open to the judge to withhold a definite answer (unless the verdict is decided by another branch: if there was no trade secret, there is no need to resolve whether the means used to obtain it were improper). It is in the passage from factors to issues that *judgements* are made and open texture is resolved. After this, resolution of the question is mechanical. So if we look below the issue level to the factor hierarchy presented in [3] we can see that, for example, *Info-Valuable* is to be resolved for the plaintiff on the basis of considering the uniqueness of the product and the degree of competitive advantage afforded. But these are not sharp concepts: use of the information may have afforded much or little completive advantage, and there may be several other more or less similar products. What is needed is a *judgement* as to whether the value of the information was *sufficient* for it to be *considered* a trade secret. It is the declaration that it is sufficient that enables the move from the extent to which plaintiff is (or is not) favoured, which is a matter of degree, to the sharp concepts needed to drive the logical model in the next stage of reasoning.

Whether or not the factors sufficiently favour the plaintiff can be decided by looking at precedent cases, or in the absence of suitable precedents, values (cf [13]): value preferences may themselves be justified by precedents as in [12] or by commentaries. Issues are considered *true* (1) if they are resolved for the plaintiff and *false* (0) otherwise.

2.3. Factors

Factors are perhaps the most studied aspect of arguing with legal cases. They are the main topic of [3], provided with a translation into rules and priorities in [25] and are the basis of the reasoning with precedents in [12] and the formalisation of such reasoning in [22]. In the *Restatement of Torts* (quoted in [7]) we find (*italics ours*):

> *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.

Note that the language here is very much non-Boolean: we consider the *extent*, *value*, *amount*, *ease or difficulty*, none of which seem likely to be answered with a "yes" or a "no". Despite this, factors in CATO (and subsequent logical treatments of the reasoning

it uses) have taken factors to be simply present or absent. Thus, for example, in [3] and subsequent work based on it, the first of the factors in the extract is represented as *Info-Known-To-Competitors*, and is said by the analyst to be present or absent in a given case, with no mention or consideration of extent.

There are two consequences of this: first the degree of doubt is hidden from view, and second that the resolution of open texture is done by the analyst with respect to factors, rather than by the judge with respect to issues. Thus while what should be determined is whether (in the view of the *judge*), taking all the factors into consideration, the plaintiff is *sufficiently* favoured by these factors for the information to be considered a trade secret, what is in fact determined is, for each feature *separately*, whether the analyst considers the extent to be sufficient for the factor to be ascribed to the case. Very often a case turns entirely on whether a factor is properly ascribed to a case or not (see, e.g., the discussion of *capture* with respect to the well known property law case of *Pierson v Post* in [11], or the discussion of *reverse engineerability* in [2]). Since this can be where the decision is effectively made, it should not be hidden in the analysis.

Thus we suggest that factors should not be seen as present or absent, but rather as favouring the plaintiff or the defendant to particular extents. One way of doing this would be to follow [18] by ascribing a weight to the factor, representing the extent to which the factor favours the plaintiff. Thus a factor can, for example, be associated with a number ranging from 1 (factor fully favours the plaintiff) to minus 1 (factor fully favours the defendant). Note that these numbers are *not* probabilities or degrees of belief but rather indicators of legal significance: hence Bayesian reasoning is not appropriate here.

In CATO, factors are arranged in a hierarchy with twenty-six *base level factors* as leaves contributing to the presence or absence of their parents, called *abstract factors*. A child is seen as a reason for the presence or absence of its parent, and the abstract factors are also taken as present or absent. Their role in CATO is that children of the same parent may be substituted for, or used to cancel, one another when emphasising or downplaying distinctions between cases with different base level factors. At the root of the hierarchy is an issue. There may be several layers of abstract factors, before the issue is reached. In CATO there are five root nodes, corresponding to the leaf issues of the logical model of IBP. Since CATO is not predicting a verdict, there is no need to tie these issues together with the logical model. As with the logical model, the abstract factors could be eliminated by unfolding into the base level factors, but they, like non-leaf issues, are useful for purposes of exposition and presentation of the argument. Thus, for example, in the case of the factors in the quotation from the *Restatement* given above, items 1 and 2 are grouped together under the abstract factor *InformationKnown*.

Factors are usually described as stereotypical patterns of fact which have legal significance. Here their legal significance is cashed out by seeing them as favouring the plaintiff (or defendant) *to some particular extent* (rather than, as in CATO and approaches based on it, favouring the plaintiff (or defendant) *simpliciter*). We now need to turn to the relation between factors and facts. It is here that the boundary between the world and the law is crossed. Recent work such as [4], and [8] has looked to dimensions, originally used in [27] and applied to law in HYPO [5], to provide the required bridge.

2.4. Dimensions

In HYPO there are a number of dimensions associated with the Trade Secrets domain. These dimensions are invoked on the basis of the facts of the cases, and are a range,

stretching from an extreme pro-plaintiff position to an extreme pro-defendant position. Thus HYPO has a dimension *Security-Measures* which has *none* as the extreme pro-defendant position and then steps through a series of more rigorous measures until the extreme pro-plaintiff position is reached. Another dimension is *Disclosures-to-Outsiders*, which ranges from the extreme pro-plaintiff position of *none* through increasing numbers of disclosures to the extreme pro-defendant point, where the information is in the public domain. Thus dimensions provide a way of determining what are the *relevant* features of the case. The set of dimensions circumscribe what can be considered a relevant aspect of the case, and the dimensions which are applicable on the facts are the aspects which are relevant to the case under consideration. Moreover since dimensions range from an extreme pro-plaintiff point to an extreme pro-defendant point they can be used to indicate the degree to which the plaintiff is favoured on that aspect. However, they do not determine the legal *significance* of the fact pattern: unless the facts put the case at one of the extremes the dimension cannot be said to favour either party. In order to determine the significance, some point or points on the dimension must be determined as cross over points, where the dimension ceases altogether to favour the plaintiff and where it begins to favour (increasingly) the defendant. The crossover may be a single point, or there may be a neutral range.

The factors in CATO can be related to the dimensions of HYPO. For example in CATO there are factors *No-Security-Measures* and *Security-Measures*. These factors divide the *SecurityMeasures* dimension of HYPO into the extreme pro-defendant point of *No-Security-Measures* and all points beyond this (*Security-Measure*s), which suggests that any security measures at all are a point in favour of the plaintiff. Corresponding to *Disclosures-to-Outsiders* we have the factors *Secrets-Disclosed-Outsiders* and *DisclosureInPublicForum*. Here the pro-plaintiff factor is not explicit: the advantage to the plaintiff of no disclosures is represented by the absence of the other two factors, which represent factors favourable to the defendant (one stronger than the other, although it is unclear in [3] if this strength manifests itself in resolving issues or is used only in emphasising distinctions).

Generally we can see the base level factors as dividing the dimension into sub-ranges. Sometimes the range will be a single point as with *No-Security-Measures*, or a quite substantial range as with *Secrets-Disclosed-Outsiders*. Thus the dimension point at which the case falls according to the facts will determine the relevance of some of the facts by allowing the case to be placed on the dimension and the legal significance of those facts by determining the base-level factor, which indicates the extent to which the plaintiff (or defendant) is favoured. Thus the analysis must determine: the identifiable points and ranges on the dimension (i.e the base-level factors); the points at which the cross over from pro-plaintiff to pro-defendant occur; and, for each sub-range within the dimension, the extent to which the plaintiff (defendant) is favoured by that sub-range. For example we can identify a dimension for closeness of pursuit with the following (increasingly pro-defendant) points, applicable to the much discussed *wild animal* cases stemming from *Pierson v Post* ([13], [11] and [10]).

> **ClosenessOfPursuit**: PhysicalPossession, MortalWounding, CertainCapture, HotPursuit, Chasing, StartedPursuit, None.

This gives a set of points, but where are we to draw the cross over point? Some authorities (e.g. Justinian) have argued that physical possession is necessary, others that *Mor-*

talWounding or *CertainCapture* are enough and Post's counsel, Livingston, argued that *HotPursuit* should suffice. It is clear therefore that all of these favour the plaintiff to some extent, and should be accorded a number greater than zero. In contrast no one has argued that the remaining three points favour the plaintiff, and so the factors representing such distant pursuit may be considered pro-defendant. The cross over thus occurs between *HotPursuit* and *Chasing*. We must therefore associate each of the four pro-plaintiff factors with a number greater than 0. The number assigned is used when we need to resolve the issue of whether the pursuit was *sufficiently close* to give possession to the pursuer. At this point the court (as it did in the case itself) may follow Justinian and require that the factor favour the plaintiff to degree 1, or it may agree with Livingston that the lesser degree of favour afforded by *HotPursuit* is enough or it may draw its line somewhere in between. The real significance of the number assigned by the analyst only emerges, however, when the factor is combined with other factors to yield values for abstract factors and issues. There are two quite different decisions here: the *analyst* chooses the cross over point from pro-plaintiff to pro-defendant but the *court* decides which party is sufficiently favoured to resolve the issue in their favour.

2.5. Legal Facts of a Case

In the previous section we saw how the use of dimension points can establish the legal significance - the degree to which they favour the plaintiff (or defendant) - of the case facts. But the dimension points have a second role: they are themselves the agreed "legal" facts of the case. Once we cross into the realm of law, the facts are taken as established, certain and indisputable. In many legal processes, including UK and US Appeals, the facts are decided by the court of first instance, and cannot be revisited at the appeal and subsequent stages. Thus, for example, in US Supreme Court cases, which are the study of much work in AI and Law, the facts are never in dispute. Two points emerge from this:

- The legal facts are taken as certain: although their *significance* can be expressed as a number n where $1 > n > -1$, the allocation to a specific range on the dimension is either 1 or 0, true or false, and it is the allocation to this range that is the legal fact.
- The potential facts are circumscribed, drawn from a limited pool. There must be a dimension to which they relate, and they must occupy a particular point or range on that dimension. This determines what can be considered as legally relevant.

2.6. World Facts

In order to move from the various degrees of certainty based on the evidence to the definite location of the case of the various dimensions, we make use of *proof standards*. The ranges on dimensions have a third role: assignments to them are the conclusions of the reasoning about the evidence presented in a case. These conclusions can be doubted to a greater or lesser extent, but the application of the appropriate proof standard ([19], [20]) enables them to be treated as definitively true (or false) for the purpose of the case. Proof standards represent the confidence in a proposition that must be attained if the court is to accept that proposition. Best known is *beyond reasonable doubt* which is a relatively high standard of proof applicable in criminal cases. Lesser standards are used

in other circumstances. The lower standard of proof in a civil case explains how a person (e.g. O.J. Simpson) can be acquitted in a criminal court and found guilty in a civil court.

Thus the factual reasoning establishes, on the basis of the evidence, a degree of belief or confidence in a statement that the case should be assigned to particular range on a dimension, and then this is accepted or rejected as a legal fact according as to whether or not it meets the standard required by the particular procedural context.

The role of ranges on dimension is thus crucial, in that they provide the bridge between world facts and legal facts. On one bank they are uncertain, as with any proposition about the world. If, however they gain entry onto the bridge by meeting the proof standard, they are taken as certain and can no longer be questioned. When they leave the bridge and enter the realm of law, they become associated with a different number, representing their legal significance, the degree to which they favour a particular party. It is very important to be aware that the passage from the world to law renders all facts equal: the degree to which they favour their party is *entirely independent* of the degree of confidence with which they emerged from the factual reasoning. Provided they cleared the bar provided by the appropriate proof standard, it matters not whether the clearance was marginal or comfortable.

Below the dimension ranges intermediate world facts are established, as part of the argument for allocation to that range. These intermediate facts will form a chain between the evidence and the dimension. These intermediate facts will have a degree of belief or certainty. We may choose to represent this as a number n with $1 > n > 0$ or with $1 > n > -1$ depending on whether we want to separate out arguments for the negations of claims, or to represent belief and disbelief on a single scale.

2.7. Evidence

At the leaves of the tree of reasoning will be the evidence. The process may appear to be akin to forward chaining: a lot of testimony is presented and the facts of the case are deduced from this body of information. In fact, however, the process is far more structured and more akin to backward chaining. Evidence is not random, but elicited by the counsels for the plaintiff and the defendant in order to establish particular points on the various dimensions. The desired points will in turn be dictated by the need to resolve the issues in favour of their clients. Thus in a Trade Secrets case counsel for the defendant will consider which issues the defendant will attempt to establish in his favour and call evidence accordingly. For example there is a choice whether to argue that the information was not a trade secret or that it was not misappropriated, and different factors relate to these different issues. There is no point in presenting evidence about security measures if the defence turns on the absence of a confidential relationship and blameless conduct by the defendant.

The evidence itself may take many forms, the testimony of eyewitnesses, expert testimony, physical evidence and perhaps even video evidence (as in *Popov v Hayashi* [10]). The evidence itself is certain: that Post testified that he was within 100 yards of the fox is a matter of record and a transcript of the testimony is available in case of doubt. Whether Post is believed or not is a different matter, especially if other witnesses testify that he was a good furlong behind the fox. The methods used to draw inferences from the evidence to beliefs as to what was the case are as many and various as the evidence presented. Witnesses may be more or less personally credible and may have

Table 1. Summary of Statement Types

Statement Type		"Truth" Value	Justification
Verdict	V	0, 1	Entailed by LI
Intermediate Issue	II	0,1	Entailed by LI
Issue	LI	0,1	$af \in AF$ Sufficiently favours Plaintiff
Abstract Factor	AF	-1 1	Deduced from BF
Base Level Factor	BF	-1 1	Mapped from DR
Legal Fact	DR	0,1	$if \in IF$ meets Proof Standard
Intermediate Fact	IF	0 ... 1	Derived from E
Evidence	E	1	Given

been better or worse placed to view the events. Expert witnesses will disagree. Statistics must be interpreted. In principle any form of standard reasoning might be called into play here. It is, of course, not specifically legal reasoning and very often the evidence is not assessed by judges but by juries composed of lay people who are supposed to be as good as anyone in deciding what is true on the basis of testimony. In AI and Law reasoning with evidence has been considered by several projects, including the hybrid (narrative and argumentation) approach of Bex [14], use of probability and Bayesian reasoning (e.g [23]) and combinations of these (e.g. [30]). Whatever the form of reasoning employed, the movement is the same. Testimony is presented, intermediate facts are derived with varying degrees of certainty, to establish a degree of belief in the particular location of the case on various dimensions, whereupon application of the appropriate standard of proof admits (or excludes) the claim as a legally accepted fact.

In summary as we progress from evidence to verdict, we pass through a number of different stages, each associated with their own particular input, output, and type of statement. The various propositions have characteristic values. These values may not be strictly *truth* values, but include the obtaining of felicity conditions, the extent to which the plaintiff is favoured, degrees of belief, etc, as appropriate to the statement type. These are shown in Table 2.

3. Representation for Systems Design

In [2] it was suggested that Abstract Dialectical Frameworks (ADF) [16] could be adapted as an effective design tool for legal case based systems, playing a role akin to that played by Entity Relationship Diagrams in the design of database systems. This is also true of what has been presented in this paper. ADFs provide a very natural tool for differentiating between statement types, by charactering statements not as one homogenous type, but as 4-tuples distinguishing the various types. Properties of the statements are:

- Statement type (as in the first column of Table 1).
- Type of parents (drawn from the first column of Table 1, except that the parent of *verdict* will be *none*).
- Type of children (drawn from the first column of Table 1, except that the children of *evidence* will be *none*).
- Value (as in the third column of Table 2).

We can also incorporate some key parameters in the acceptance conditions. For example the acceptance conditions for *issues* will require a threshold indicating the degree to which the plaintiff (defendant) must be favoured for the issue to be resolved in their favour, and the acceptance conditions of *legal facts* will need to incorporate the proof standard that is applicable. We could, if desired, impose further conditions on the form of the ADF required to give an acceptable design. For example, while all designs have, in our experience [2], naturally produced cycle free ADFs, this could be made a requirement for an acceptable design, since the presence of cycles would present problems at the implementation stage. We might also wish to constrain the design in other ways: for example [1] commends the use of *two-regular ADFs* in which every non-leaf node has exactly two children. This has some software engineering advantages, but requires the use of additional intermediate components. It does not affect the logic of the ADF, and so may be left to the taste and predilections of the user of the design tool.

Additionally we will need a small ontology recording dimensions, the recognised ranges on the dimensions, the cross overs and the degree to which these points are said to favour the plaintiff. The size is tightly constrained by the use of dimensions to indicate relevance, and so does not fall foul of the problems identified in [6], where producing the ontology seems a dauntingly unlimited task. The problems noted there do, however, come to the fore when trying to design the factual reasoning from evidence to the legal facts. This is because the facts - and the appropriate styles of reasoning - are so disparate. But there is no necessity for a system to include this level of detail. The layers provide a number of different entry and exit points. HYPO for instance enters at dimensions, and CATO at factors and the British Nationality Act program at issues, and all three provide useful systems. Similarly HYPO and CATO exit below the level of a verdict. Bex in [14] enters at evidence, but exits below factors. While it is important to scope a system, there is no need to address the whole problem in a single system.

4. Concluding Remarks

In this paper we have attempted to give a complete picture of reasoning with legal cases from evidence to a verdict. In doing so we are able to locate previous work, which tends to address parts of the process, within the larger whole, and so bring together the different contributions made by different approaches. In particular we have identified the different types of statements which are the subject of different stages, and to explain why different types of reasoning are appropriate at different stages: why we can sometimes fruitfully use Bayes when reasoning from evidence to legal facts, but not when reasoning from factors to issues, for example. We have also shown where reasoning should be seen as Boolean and where various forms of uncertainty and degree are introduced and eliminated. Finally we have drawn attention to the pivotal role of dimensions in bridging the gap between the realm of the world and the realm of law.

References

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

[2] L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artificial Intelligence and Law*, 24(1):1–50, 2016.

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

[4] M. Araszkiewicz, A. Łopatkiewicz, and A. Zienkiewicz. Factor-based parent plan support system. In *Proceedings of the 14th International Conference on AI and Law*, pages 171–175, 2013.

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

[6] K. D. Ashley. Ontological requirements for analogical, teleological, and hypothetical legal reasoning. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, pages 1–10, 2009.

[7] K. Atkinson and T. Bench-Capon. Legal case-based reasoning as practical reasoning. *Artif. Intell. Law*, 13(1):93–131, 2005.

[8] K. Atkinson, T. Bench-Capon, H. Prakken, and A. Wyner. Argumentation schemes for reasoning about factors with dimensions. In *Proceedings of JURIX 2013*, pages 39–48. IOS Press, 2013.

[9] J. L. Austin. *How to do things with words*. Oxford university press, 1975.

[10] T. Bench-Capon. Representing Popov v Hayashi with dimensions and factors. *Artif. Intell. Law*, 20(1):15–35, 2012.

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

[13] D. Berman and C. Hafner. Representing teleological structure in case-based legal reasoning: The missing link. In *Proceedings of the 4th International Conference on AI and Law*, pages 50–59, 1993.

[14] F. Bex. *Arguments, stories and criminal evidence: A formal hybrid theory*. Springer, 2011.

[15] 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, 1991.

[16] G. Brewka, H. Strass, S. Ellmauthaler, J. P. Wallner, and S. Woltran. Abstract dialectical frameworks revisited. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 2013.

[17] S. Brüninghaus and K. Ashley. Predicting the outcome of case-based legal arguments. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 233–242, 2003.

[18] 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.

[19] A. M. Farley and K. Freeman. Burden of proof in legal argumentation. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 156–164. ACM, 1995.

[20] T. F. Gordon and D. Walton. Proof burdens and standards. In *Argumentation in artificial intelligence*, pages 239–258. Springer, 2009.

[21] H. L. A. Hart. *The concept of law*. Oxford University Press, 2012.

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

[23] J. Keppens. Towards qualitative approaches to bayesian evidential reasoning. In *Proceedings of the 11th international conference on artificial intelligence and law*, pages 17–25. ACM, 2007.

[24] L. Lindahl. Deduction and justification in the law. the role of legal terms and concepts. *Ratio Juris*, 17(2):182–202, 2004.

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

[26] A. Rigoni. An improved factor based approach to precedential constraint. *Artif. Intell. Law*, 23(2):133–160, 2015.

[27] E. L. Rissland. The ubiquitous dialectic. In *Proceedings of the Sixth European Conference on Artificial Intelligence*, pages 367–372, 1984.

[28] M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. The British Nationality Act as a logic program. *Communications of the ACM*, 29(5):370–386, 1986.

[29] D. Skalak and E. Rissland. Arguments and cases. *Artif. Intell. Law*, 1(1):3–44, 1992.

[30] S. T. Timmer, J. C. Meyer, H. Prakken, S. Renooij, and B. Verheij. A structure-guided approach to capturing bayesian reasoning about legal evidence in argumentation. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 109–118, 2015.

[31] F. Waismann. The principles of linguistic philosophy. 1965.