

HYPO'S Legacy: Introduction to the Virtual Special Issue

T.J.M. Bench-Capon

the date of receipt and acceptance should be inserted later

Abstract This paper is an introduction to a virtual special issue of *AI and Law* exploring the legacy of the influential HYPO system of Rissland and Ashley [116] and [19]. The papers included are:

- Arguments and cases: An inevitable intertwining [134].
- BankXX: Supporting legal arguments through heuristic retrieval [124].
- Modelling reasoning with precedents in a formal dialogue Game [108].
- A note on dimensions and factors [117].
- An empirical investigation of reasoning with legal cases through theory construction and application [66]
- Automatically classifying case texts and predicting outcomes [21].
- A factor-based definition of precedential constraint [91].
- An improved factor based approach to precedential constraint [110].

After describing HYPO, in this introduction to the special issue I look at various aspects of its influence on AI and Law: the developments led by Rissland at Amherst; the developments led by Ashley in Pittsburgh; the expression of these ideas in terms of rule based systems, and their subsequent formalisation; value based theories, which were inspired by a critique of HYPO; and contemporary approaches which revive the idea of dimensions.

1 Introduction

It is now thirty years since HYPO came to the attention of the international AI and Law Community at the very first ICAIL at which [116] and [23] were presented. Although Edwina Rissland had published previously on reasoning with hypotheticals in [111], [126] and [113] and its application to law in [112], it was the PhD project with her student Kevin Ashley which culminated in [19] which

T.J.M. Bench-Capon
Department of Computer Science
University of Liverpool
Liverpool, UK
E-mail: tbc@csc.liv.ac.uk

is generally regarded as the HYPO system (first referred to as such in [127]). Over the past 30+ years HYPO has been without doubt the most influential AI and Law project, setting the agenda for reasoning with legal cases, arguing with legal cases, rule based approaches to reasoning with cases and formalisations of precedential reasoning. The papers in this virtual special issue, which are drawn from all three decades, represent a range of significant papers which have built upon and developed ideas from the HYPO system. In this introduction I will use these papers to illustrate HYPO's considerable legacy. There are, of course, many other papers, a number of which were published in other venues which are needed to tell the whole story, but the selection in this issue is fairly representative of the way thinking in AI and Law has developed. A number of other papers will be referred to in this introduction, so that the complete story can be told.

Section 2 describes HYPO and its key features, in particular its representation of cases as *dimensions*, and its conception of the form of a *legal* argument. Section 3 describes some direct descendants of HYPO, led by Rissland at Amherst: CABARET [134]*¹, which combined rule-based and case-based reasoning; BankXX [124]* which envisaged constructing legal arguments by heuristic search over a rich network of argument pieces; and SPIRE [119], which used HYPO derived techniques to drive a system to retrieve legal cases. Section 4 considers the developments from HYPO led by Ashley following his move to Pittsburgh. The most influential is CATO, carried out with his student Vincent Alevan, [13]. CATO was never the topic of an article in this journal, and is most completely reported in [10] and [11]. CATO is best known for its move from dimensions to *factors*, and its organisation of these factors into a *factor hierarchy*. A subsequent project with another of Ashley's students, Stefanie Brüninghaus, led to Issue Based Prediction (IBP) [21]*. During the 1980s and the early 1990s, case-based and rule-based approaches to legal reasoning were seen as distinct. Some saw them as potentially complimentary as in CABARET, but others, especially in Europe with its logic programming and civil law traditions, saw them as opposed to one another. An important paper by Prakken and Sartor [108]*, however, demonstrated how a set of precedents could be represented as a set of rules and priorities between them. This enabled formal accounts of reasoning with cases to be developed [91]* and [110]*. Representation as rules also permitted the use of argumentation schemes to support reasoning with cases: [144], [39] and [109]. The move to rules and the formalisation of the reasoning is described in section 5. Another direction was initiated by [56], which led to preferences being explained in terms of purpose and values. This idea was formalised as theory construction in [53], and explored empirically in [66]*. These developments are explored in section 6. Although HYPO is the *fons et origo* of all this work, mostly it has been based on factors, as presented in CATO. There has, however, been interest in reviving dimensions to allow more nuance to be captured. The differences between dimensions and factors are discussed in [117]*. Although the return to dimensions was argued for in [50], it is only more recently that work on this has intensified: [109], [17] and [6]. This work is discussed in section 7, and some concluding remarks are made and open questions identified in section 8.

¹ The papers in this virtual special issue are indicated with an “*”.

2 HYPO

Although the ideas relating to HYPO were around earlier, (a panel was held at IJCAI 1985 [114], with panellists including both Rissland and Ashley, together with Michael Dyer, Anne Gardner, Thorne McCarty and Donald Waterman), details of HYPO itself began to emerge at the first ICAIL in 1987 [116]. By 1989 Ashley had moved to Pittsburgh and so the discussion of HYPO at the second ICAIL was his sole authored [18], and his 1988 PhD thesis appeared in book form as [19], providing the most detailed account of the project. HYPO operated in the domain of US Trade Secrets Law. Essentially the law protected trade secrets against misappropriation. To be a trade secret the information needed to have value, and to be regarded as a secret by the plaintiff, shown by making reasonable efforts to maintain secrecy. Misappropriation took the form either of using information knowing it to be confidential, or acquiring the information by dubious means, such as deception or bribery. Defences typically turned on the availability of the information through legitimate sources, or independent development of the information.

The most important ideas behind HYPO were dimensions, and a particular conception of argumentation, namely three-ply argument. These will be discussed in turn in the following subsections.

2.1 Dimensions

A major difficulty when thinking about how to reason with legal cases is that cases which are held to be similar appear to be very different when we consider their facts. Thus in a famous line of cases considered in [94], [44] and [128] a collection of very different things are seen to play the same role in the various cases². These include a scaffold, an elevator, a bottle of aerated water and a coffee urn which were all held to be relevantly similar to an automobile for the purposes of the particular legal point under consideration. Similarly the series of property cases much studied in AI and Law consider a fox, wild ducks, a shoal of pilchards and a whale all to be relevantly similar to the baseball disputed in *Popov v Hayashi* (see [24]). Thus deciding whether two legal cases are relevantly similar does not seem to be a matter of considering facts in a simple manner using everyday notions of similarity.

For this reason HYPO introduced the notion of *dimensions* to match and compare cases. These are aspects of the case, ascribed on the basis of facts, which are relevant to the legal issues being considered. Thus what linked the scaffold, elevator, bottle of aerated water and coffee urn was that all were dangerous, but not visibly so. What links the animals and the baseball is that they were being pursued, and the relevant aspects include their value, and their connection to the land they were pursued on [43].

² The need to consider disparate types of thing as if they were the same sort of thing is pervasive in analytic philosophy. Austin wrote in [30] "But then we have to ask, of course, what this class comprises. We are given, as examples, 'familiar objects'—chairs, tables, pictures, books, flowers, pens, cigarettes; the expression 'material thing' is not here (or anywhere else in Ayer's text) further defined. But does the ordinary man believe that what he perceives is (always) something like furniture, or like these other 'familiar objects' - moderate-sized specimens of dry goods?"

Dimensions are aspects of a case which may or may not be applicable. If applicable, the dimension represents a vector, taking a range of values which entirely favour one of the parties at one end and then increasingly favour the other party, until at the other end of the range the dimension entirely favours that other party. At some point the dimension will cease to favour the plaintiff and at some point (possibly the same point) it will start to favour the defendant. Thus if there are n dimensions, we have a n -dimensional space, with the plaintiff favoured at some locations and defendant favoured at others.

We can illustrate dimensions using the domain of the Automobile Exception to the 4th Amendment, discussed in [115], [22] and [7]. This is a good illustration because it can be seen as having just two dimensions, privacy and exigency, which allows the space to be pictured, as in Figure 1. Essentially the 4th Amendment provides that

“[t]he right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, ...”

The automobile exception permits search of an automobile without a warrant provided that there is probable cause, and that the search is sufficiently exigent. Exigency arises from the possibility that the delay in seeking a warrant might compromise the collection of evidence or public safety. If we consider the dimension of exigency, a completely immovable item would represent one end of the dimension and an automobile in transit on a highway the other. If we arrange our precedents along this dimension we will observe that at one extreme, where there is little exigency, all the cases will require a warrant, and at the other, where the matter is pressing, none of the cases will require a warrant. In between, some cases will need a warrant and some will not, suggesting that other relevant dimensions are in play. The situation is represented pictorially by the vertical grey lines in Figure 1³. Similarly for privacy, at one end (e.g. intimate body search) reasonable expectations of privacy will be such that a warrant will always be required, whereas at the other (in plain view in a public space) there may be little or no expectation of privacy so that a warrant may never be required. This situation is shown by the horizontal lines of Figure 1. Putting the two dimensions together, the lines divide the space into 9 two-dimensional areas. In A, D and G, there is insufficient exigency to merit warrantless search, whereas in C and F the exigency is such (perhaps an immediate threat to the life of the President) that no expectations of privacy are sufficient to require a warrant. In A, B and C there are insufficient expectations of privacy to require a warrant given any degree of exigency, while in G and H the expectations of privacy are sufficiently high that no degree of exigency will permit the search. The interesting areas are E and I. In E there is a trade off between the two dimensions, so that some kind of balance must be struck [93]. In area I, unless there is another dimension to consider, a preference needs to be expressed, perhaps based on a preference between values [53]. The dark line in Figure 1 thus divides the space into areas where a warrant is required (to the left

³ A similar diagram was actually used in the decision in the Massachusetts Trade Secrets case *Sissy Harrington-McGill v. Old Mother Hubbard Dog Food, Inc.*, Middlesex Superior Court C.A. 81-4385.

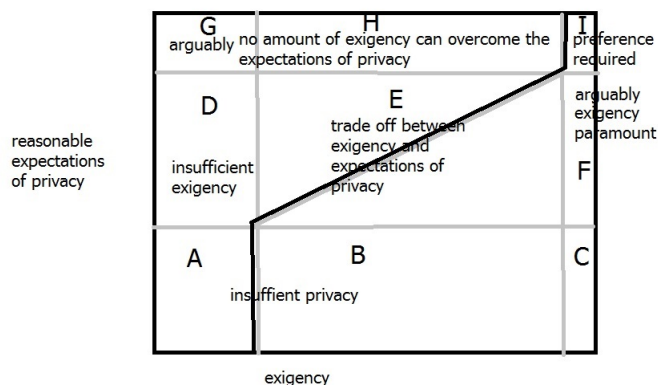


Fig. 1 Automobile Exception Using Two Dimensions

of the line) and those where it is not (to the right of the line)⁴. The line shown in Figure 1 indicates an even trade off between privacy and exigency (giving the 45° line): other trade-offs are possible, which would require a curve rather than a line. It also includes area I in the no-warrant space, indicating a preference for exigency: a different preference would send the line horizontal, excluding I. Note also that trade-off could be possible throughout the range of the dimensions, so that areas A and I are squeezed. The right hand grey line effectively represents the point at which the dimension becomes decisive (subject to preference for a second dimension which is the decisive range for the other side), what [62] call a *knock-out* factor. If both dimensions are such that the area E expands to consume the whole space, we have a situation where the dimensions are never decisive, but require to be balanced across their whole range.

If there is a third dimension the space will become a cube, with 27 partitions, and more dimensions will make visualisation very hard, although mathematicians can happily write equations relating to high dimensional spaces⁵. If we now position the precedents in this dimensional space we can see the problem as being to divide the space into the plaintiff region and the defendant region. In the two dimensional case of Figure 1, this could be achieved by the black line with the warrant cases on the left and the non-warrant cases on the right. The problem would seem to be the kind of line fitting problem which could be addressed by nearest neighbour or other statistical techniques. But there are two reasons why such techniques are not suitable here: first it is quite clear that this is not the way that people reason with cases, and so the point of the exercise, which was to model lawyer-like reasoning, would not be achieved. But more importantly there are not really sufficient cases to apply the statistical techniques, which depend heavily on having abundant examples. HYPO has 13 dimensions and many of the cases in fact only use some of these, typically no more than 4 or 5. The use of

⁴ Those familiar with machine learning will recognise the similarities with a decision boundary in supervised learning.

⁵ For example $x^2 + y^2 = r^2$ is a circle of radius r , $x^2 + y^2 + z^2 = r^2$ is a sphere with radius r , $x^2 + y^2 + z^2 + w^2 = r^2$ is a four dimensional “4-sphere” sphere with radius r , and so on.

line fitting techniques would require hundreds or even thousands of cases (for a discussion of the required number of cases see [4]). To obtain, let alone perform a dimensional analysis on, that many cases was far beyond what was sensible to ask from a nineteen-eighties PhD project. HYPO has 33 cases in its case base [19]. It remains a task few would undertake, even if cases were available in this quantity⁶.

HYPO models argumentation, rather than using a purely statistical approach. Given a set of precedents located in the n -dimensional space, and a new case located in that space, the idea is to find arguments, based on the precedents, for and against a finding for the plaintiff (or defendant). Given a plaintiff precedent P_p and a defendant precedent P_d and a current case C_c , the plaintiff's task is to argue that C_c is on the same side of the line (or its n -dimensional equivalent) as P_p rather than P_d . For example, in Figure 1, if C_c is South East of P_p , that is a strong argument for the plaintiff: if it is South West (or North East) we have arguments for and against the plaintiff, and if it is North West, it is of no help at all to the plaintiff.

In HYPO the arguments were deployed in a Three-Ply argumentation structure, as described in section 2.2.

2.2 Three-Ply Arguments

HYPO deploys its arguments within a three-ply structure. This structure is also found in the presentation of oral arguments in the US Supreme Court [2], but in the Supreme Court the plies take the form of a dialogue between counsel and the Justices, whereas HYPO arguments are presented uninterruptedly (as if by the counsel), rather than in a dialogue. The three plies are:

1. **Cite a Case:** One side (say the plaintiff) cites a case found for that side. That case should be as similar (*on-point*) as possible to the current case, and the suggestion is that the decision in the precedent should be applied to the current case (following the principle of *stare decisis*, which means *let the decision stand*).
2. **Response:** The defendant responds to the plaintiff by distinguishing (that is, pointing to significant differences which mean that the precedent should not be followed), and by citing counter examples, precedent cases found for the defendant that are at least as on-point as the case cited by the plaintiff.
3. **Rebuttal:** The plaintiff now attempts to rebut the contentions made by the defendant in the second ply: distinguishing the counter examples, emphasising any strengths and attempting to show that any weaknesses are not fatal.

Each of these plies will be explained below.

2.2.1 Citing a Case

HYPO measures similarity between cases in terms of four metrics: on pointness (the degree of overlap in terms of dimensions), outcome, the magnitudes of shared dimensions and potential relevance as a near miss. To assess on pointness, cases are

⁶ Where large quantities of cases are available, they tend to be subjected to machine learning techniques rather than manual analysis. A recent example is [9], which will be discussed in more detail later (see footnote 12).

(partially) ordered in a *Claims Lattice*, as shown in Figure 2. The lattice includes every case which shares at least one dimension with the case being considered. Figure 2 shows the Claims Lattices for *USM Corp v Marson Fastener Corp*⁷. Note that a precedent is considered more on point than another precedent only if the dimensions of the second precedent form a proper subset of the dimensions of the first. Thus *Space Aero v Darling*, which shares three dimensions with *USM*, is not considered more on point than *Automated Systems*, which shares only two, because the dimensions of *Automated Systems* are not a proper subset of *Space Aero v Darling* and so they are on different branches of the lattice. The most on point case, decided for the desired party, and which does not have a more on point case for the other party is the one to cite. Of the three first level cases in Figure 2, only *Space Aero v Darling* was found for the plaintiff, and so will be cited in the first ply. If there are equally on point cases, several arguments, based on the different cases, could be made in the first ply.

2.2.2 Response

The defendant must now provide a response. First, an attempt to distinguish *Space Aero* should be made. Distinctions can be made by showing that the precedent is stronger for the plaintiff, or by showing that the current case is stronger for the defendant. One way in which this can be done is in terms of different dimensions: if there is a dimension present in one case but not the other, that can provide a reason for saying that the precedent does not govern the current case, although it might still do so if there are also reasons for finding that this distinction is not sufficient. This issue was explored in more depth in CATO, as discussed in section 4.1. To be a good distinction the different dimension should represent a strength (weakness) for the plaintiff if it is in the precedent (current) case, and a strength (weakness) for the defendant if it is the current (precedent) case. The differences between these two ways of distinguishing is discussed in [39]. Thus in *USM*, that secrets were disclosed to outsiders in *USM* but not *Space Aero* could be used as a distinction, since it represents a weakness for *USM*. A second way in which a precedent can be distinguished is on a shared dimension. Thus if the information gave a larger competitive advantage to the defendant in *Space Aero* than *USM*, that could be a point of distinction.

Having made the distinctions between the plaintiff's precedent and the current case, the defendant may now cite counter examples of his own. Both *Automated Systems* and *Crown Industries* would be possible counter examples since they are as on point as *Space Aero* and found for the defendant.

2.2.3 Rebuttal

The rebuttal stage gives the plaintiff the opportunity to answer the points made by the defendant in the response. Three types of answer are possible:

- *Distinguish the counter examples.* This can be done in the same ways as the defendant distinguished the plaintiff's precedent in the response.

⁷ 379 Mass. 90 (1979). This case which concerned the manufacture of blind rivets, turned largely on whether the security measures taken were adequate.

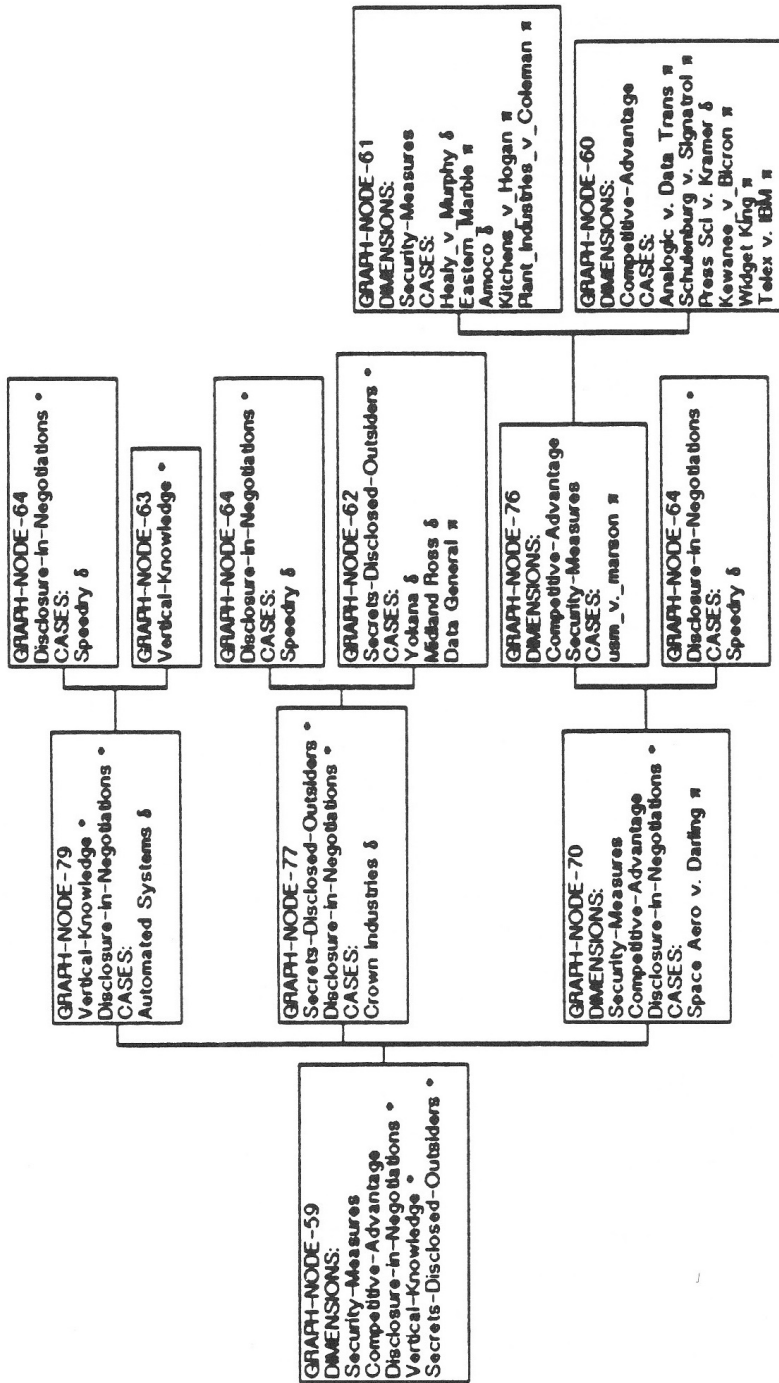


Fig. 2 Claims Lattice for USM, Figure 8.2 in [19]

- *Emphasise strengths.* If the current case is more favourable to the plaintiff on some dimensions than a precedent which already favours the plaintiff, that is a strength worth emphasising.
- *Show weaknesses not fatal.* If the defendant has distinguished a case by showing that the plaintiff is weaker on some shared dimension than a precedent found for the defendant, this can be shown to be not a fatal weakness by finding another precedent which is weak or weaker on that dimension but which was found for the plaintiff. This shows that we are in an area of that dimension which has mixed outcomes, suggesting that it is some other dimension that is decisive.

An important idea in HYPO (the name of the PhD project derived from Rissland's earlier work on hypothetical reasoning) is that if the current case falls between precedents with different outcomes, hypothetical cases can be used, varying the position on the dimension, or by adding or removing a dimension, so that intuitions can suggest exactly where the line should be drawn, and on which side the current case should fall. This aspect of HYPO was perhaps less exploited in the HYPO system in its final form than might have been hoped.

Thus HYPO has a structure within which to deploy the extensive range of arguments made possible by applying its dimensional analysis to the cases. In the next section I will discuss the thirteen dimensions used by HYPO in [19].

2.3 Dimensions in HYPO

As mentioned before, HYPO uses thirteen dimensions. But these are not homogeneous. The paradigmatic dimension is continuous, running from the extreme plaintiff to the extreme defendant location. But some of the dimensions in HYPO are not continuous: rather their range comprises a set of enumerable points. In other examples the dimension degrades to just two points: the extreme plaintiff and the extreme defendant point, and no intermediate points are recognised. Finally the dimension may degrade to a single point: for example if the defendant paid a bribe that is a strong point for the plaintiff, but if no bribe was paid that does not help the defendant's case, rather it means that the dimension does not apply to the case. HYPO's dimensions and their types are listed below:

1. *Competitive Advantage Gained.* This dimension combines the development costs in terms of time and in terms of money. It is not entirely clear how these are combined, but the dimension can be regarded as **continuous**. It might have proved convenient to split into two dimensions: time savings and money savings.
2. *Vertical Knowledge:* Vertical knowledge is contrasted with technical knowledge, and vertical knowledge favours the defendant. This is a **binary** dimension,
3. *Secrets Voluntarily Disclosed:* This is simply the number of people to whom there was voluntary disclosure. It is **continuous**.
4. *Discloses Subject to Restriction:* This is the percentage of disclosures which are subject to restriction. As such it is potentially continuous, but in [19] only 0% or 100% are recognised and so it is effectively **binary**.
5. *Agreement Supported by Consideration:* This refers to a (confidentiality or non-disclosure) agreement between the plaintiff and the defendant, and is a **binary**

dimension: the amount of the consideration is not thought important in [19], as is usual in law. *Peppercorn rents* are an example where the consideration is purely nominal.

6. *Common Employee Paid to Change Employers*: Although this again could in principle be considered continuous, no account is taken of the number of employees concerned, nor of the amount paid, and so it is better regarded as **binary**. In fact, since the dimension does not really apply if there are no common employees, it is effectively **unary**.
7. *Exists Express Noncompetition Agreement*: This refers to an agreement between the plaintiff and former employees. It is **binary**.
8. *Common Employee Transferred Product Tools*: This is best regarded as **unary**: the dimension does not apply if there are no common employees, or they did not bring tools.
9. *Non-Disclosure Agreement Re Defendant Access*: This is an agreement between plaintiff and defendant, and is **binary**.
10. *Common Employee Sole Developer*: If the sole developer transferred employer, it may be considered reasonable that he takes his knowledge with him. This is a **binary**, or perhaps even unary, dimension.
11. *Non-disclosure Agreement Specific*: This dimension only applies if there is a non-disclosure agreement between the plaintiff and the defendant, and no degrees of specificity are recognised, so that it is **binary**.
12. *Disclosure in Negotiations with Defendant*: Such disclosures can weaken the plaintiff's case. It is a **binary** dimension.
13. *Security Measures Adopted*: This an **enumerated** range with eight points, ranging from the extreme pro-defendant point (minimal measures) to the extreme pro-plaintiff point (employee non-disclosure agreements).

After HYPO, Ashley moved to University of Pittsburgh (specifically to the School of Law and Learning, Research and Development Centre). At Pittsburgh, Ashley developed CATO, with Vincent Aleven [10], and SMILE and Issue Based Prediction (IBP) with Stefanie Brüninghaus [21]. Meanwhile Rissland continued to develop the HYPO ideas at Amherst producing CABARET with David Skalak [134], BankXX with Skalak and M. Timur Friedman [124] and SPIRE with Jody Daniels [119]. A major difference between these two strands was their use and understanding of *factors*. As explained in [117], both groups adopted the more legal sounding term “factors” in preference to the mathematical term “dimensions”. But whereas at Amherst, factors were just a different term for HYPO’s dimensions, a related but different notion was used in Pittsburgh. The Amherst systems will be discussed in section 3 and the Pittsburgh systems, and the notion of factor as used there, in section 4.

3 Beyond HYPO

Work on reasoning with legal cases was continued at Amherst by Edwina Rissland, working with a succession of PhD students, David Skalak, M. Timur Friedman and Jody Daniels. Two of the resulting systems are described in papers in this special issue, [134] and [124], and the details of these two systems will be left to them. The domains for these systems are the *Home Office Deduction* (an allowance

against income tax when part of one's home is used as an office) for CABARET and personal bankruptcy for BankXX and SPIRE.

3.1 CABARET

CABARET was described in conference papers, including [121], [120] and [133], before being consolidated in [122] and [134] (the very first paper to appear in this journal), which focussed in particular on strategies and argument moves. Because it took as its domain Home Office Deduction, there was a basis for the law in statutes, whereas HYPO had started from case law and its consolidation in the *Restatement of Torts*. This enabled CABARET to take as its focus the interpretation of terms found in the statutes, the idea being that the meaning of such terms is clarified and refined in case law. Thus while the statutes can be represented using rules, much in the manner of [132], when “the rules run out” [72], it is necessary to use case based techniques. The process is well explained by Ron Loui in his commentary on [133] in [40]. In this way the rules are able to focus the search and partition the case base so that there are fewer irrelevant distinctions made during the response (and the case base can provide precedents that govern more cases, since differences which relate to terms not at issue are not considered)⁸. This use of high level rules to focus and drive the lower level case based reasoning was a significant improvement, also used in IBP [62] and still in use today [7].

Also CABARET developed a taxonomy of argument moves to represent patterns of actual legal argument, such as straw man and make-weight arguments, and showed how these were supported by the dimensional analysis of precedent cases. This was an important development since it enabled consideration of *strategies* for deploying arguments, an idea later developed through the use of *dialogue games*, such as The Pleadings Game [73], the game of [107], TDG [45], and PADUA [143], amongst others.

Both of the aspects introduced in CABARET represent significant developments of HYPO, allowing the dimensional analysis of cases to be employed in the context of rule interpretation, and for arguments with a more natural and familiar form to be used a strategic manner.

3.2 BankXX

BankXX, which operates in the domain of personal bankruptcy (specifically Chapter 13 of the U.S. Bankruptcy Code), was most fully reported in [124] (in this issue) and its evaluation described in [125], an earlier conference version having appeared as [123]. BankXX represents a considerable departure from HYPO and CABARET, since it uses the precedents (and other sources) to represent the domain knowledge as a highly interconnected network of building blocks which which is searched to gather *argument pieces*. The nodes in this network encompass a wide variety of ways of representing the domain knowledge including cases as collections of facts, cases as dimensionally-analyzed fact situations, cases as bundles of citations, and cases as prototypical factual scripts, as well as legal theories represented in terms of domain dimensions. Thus cases are represented in several

⁸ An alternative approach to reasoning with portions of precedents is offered by [59].

ways. In particular, in its *Domain Factor Space*, cases are represented “by a vector composed of the magnitudes of the case on each dimension that applies to it; non-applicable factors are encoded as NIL. This ... represents a case as a point in an n-dimensional space.” Arguments are then formed by performing heuristic search over the network, using evaluation functions at the domain level, the argumentation piece level, and the overall argument level. The result is a highly sophisticated system which can blend the more Boolean nature of CATO’s factor-based approach with the more value-oriented nature of HYPO’s dimension-based approach. Further, the factors are grouped according to various theories, which, as in CABARET, gives some structure to the set of all factors.

The evaluation in [125] is one of the most (if not the most) detailed examples of evaluation in AI and Law. It considers several different forms of the BankXX program, and the evaluation is conducted from several perspectives, and a number of issues relating specifically to the evaluation of programs in the domain of law are noted.

BankXX has been rather unduly neglected⁹, and has no obvious descendants in current AI and Law research. This does scant justice to the importance of the work. Construction of cases by performing heuristic search was also carried out by AGATHA [65], but the search tree was over only a collection of cases represented as bundles of CATO-style factors, rather than the highly sophisticated network of knowledge used in BankXX. Perhaps AI and Law should make more use of traditional AI techniques such as heuristic search. It is to be hoped that the appearance of BankXX in this special issue may revive interest in this work, and perhaps lead to more thorough evaluations of work in AI and Law. Evaluation remains an issue in AI and Law [68].

3.3 SPIRE

SPIRE [118], [119] and [69], which was applied to both the domain of home office deduction and the domain of personal bankruptcy, used case based reasoning to kick-start (*seed*) an information retrieval system. The idea was that in this way a large case base could be searched, but only a small number of these cases need be given the detailed analysis required for dimensional case based reasoning. From a case based reasoning view point this gives access to very much larger *corpora* of cases, while from an information retrieval point of view it enables the formulation of better queries.

The system operates in two stages (originally developed separately: the name SPIRE was not used in the earlier papers).

- First [118], [119], the analysed cases are used to retrieve documents (case decisions) that are relevant to the presented problem case, and
- Second, within those retrieved documents, passages that contain information relevant to specific case features are highlighted. Here the “case base” (more like an examples base) was a collection of excerpts; one set for each term of interest (such as the *sincerity* of the debtor). Differences with other systems in the HYPO family are that: they are not structured objects like cases, (but

⁹ As of March 2017, [124] had received only 83 citations on Google Scholar, and the evaluation in [125] a mere 19.

only snippets of text; the same set of excerpts is used for all attempts to locate passages and thus, there is no attempt to make the selection of “see” excerpts “relevant” to the case at hand. The idea was then to do something analogous for case retrieval, but this time on an individual document and to locate relevant passages within it.

These two stages were put together as SPIRE [69]. For SPIRE, the input was a fact situation and the output was a set of relevant passages in each of the most relevant cases. It would then be possible to use these passages to query the main case base.

The engine used to query the main case base was INQUERY [64], developed by the University of Massachusetts Information Retrieval group in the Center for Intelligent Information Retrieval (CIIR). The use of HYPO is in selecting the seed cases: the top two layers of the case lattice proved to be a good choice for determining the cases to use.

Several different queries were tried, including bag of words, set of words and what was termed a “sum” query which required terms from within an excerpt to be found co-occurring in the passage, and some weighting methods. Among the results was that bag of words performed better than set of words, and bag of words and sum performed better than the other query types tried.

Although an interesting use of legal case based reasoning, and with some encouraging results, SPIRE was not widely influential. This may be because Daniels went on to have a successful career in the US Army after finishing at Amherst, but in any case information retrieval is a hard field for research to make a lasting impact in because of the rapid development of technology. The tools now available for information retrieval are such as could only have been dreamed of in the mid-90s when the SPIRE research was carried out (which was before Google transformed the way we access information).

4 Doing Things with Factors

Following HYPO, Kevin Ashley moved to Pittsburgh in 1989, and continued to work on reasoning with legal cases, using many of the ideas developed in HYPO. Significantly, Ashley moved specifically to the Learning and Development Research Center. This connection with the Learning Center had a significant influence on Ashley's projects, in that the task focus became legal education. Also the position gave access to real classes of legal students for evaluation of these systems. Ashley took full advantage of this opportunity and his work at Pittsburgh is marked by its focus on empirical evaluation and the use of carefully designed experiments to produce empirical results. His student, Vincent Aleven, with whom he developed CATO (first mentioned as such in [12]), entitled his PhD *Teaching case-based argumentation through a model and examples* [10], and it made extensive use of law students in its evaluation. The key systems we will look at in this introduction, because of their influence on AI and Law, are CATO and IBP [62] and [21]. Other projects carried out by Ashley at this time included Sirocco [99] with Bruce McLaren.

4.1 CATO

The task CATO was intended to address was to support teaching law students to distinguish cases. Not every difference between a case and a precedent represents a usable distinction. If we think in two dimensions, as in Figure 1, the plaintiff “owns” the area to the south-east of the precedent. Therefore a distinction has to place the current case to the north or west of the precedent since otherwise the difference actually strengthens the case, and the current case can be decided in the same way as the precedent using an *a fortiori* argument. Even if the distinction is usable, however, it may be that there are other differences between the current case and the precedent which can be used to compensate for or cancel out the identified distinction. Such a difference enables the distinction to be *downplayed*. If a distinction cannot be downplayed it may be considered significant, and so can be emphasised. Thus the task of CATO was to teach students:

- which differences between a case and a precedent can be used as distinctions,
- which distinctions can be downplayed,
- how to downplay distinctions.

4.1.1 Factors in CATO

A major difference between HYPO and CATO was the use of *factors* instead of dimensions. In HYPO cases had been represented in terms of their facts, and the facts used to determine whether the dimension was active in a particular case and, if active, where on the dimension the case lay. Factors are also related to facts, but are attributed by the analyst on the basis of the case facts, so that factors are simply present in or absent from a case and cases can be represented as bundles of factors. Thus in CATO facts are seen only by the analyst. Another important feature of factors, is that the presence of a factor always favours the same side: a factor is *either* pro-plaintiff *or* pro-defendant. In sum factors can be considered as *stereotypical patterns of fact that favour one party or the other*. While this is something of a simplification when compared to dimensions, it does not affect the task being taught: the difference between cases can be considered in terms of differences between the factors present in the cases: the process of attributing factors is not important to this particular task.

Because HYPO and CATO both use the same domain, US Trade Secrets, direct comparison can be made between CATO’s factors and HYPO’s dimensions. CATO used 26 factors (F1-F27 - there was no F9), as shown in Table 1. The Table also shows any relation of the factors to HYPO dimensions: Dimensions are identified using the numbers in section 2.3.

All but one of the dimensions (which relates to contract law rather than trade secrets) in HYPO correspond to one or more factors. In addition there are a number of factors which are unrelated to HYPO dimensions: these mostly relate to legitimate and questionable means of discovering the information. CATO analysed more cases (148 as against 33) and so it is to be expected that new factors or dimensions will emerge from the analysis. In fact if we replace D2 with a dimension *Questionable Means*, and introduce a new dimension *Legitimate Means* we can locate most of these new factors on these dimensions. Or we could see these as a single dimension with F22 as the extreme pro-plaintiff point and F24 as the

ID	Factor	HYPO Dimension
F1	DisclosureInNegotiations (d)	Defendant end of D12
F2	BribeEmployee (p)	D6
F3	EmployeeSoleDeveloper (d)	D10
F4	AgreedNotToDisclose (p)	Plaintiff end of D9
F5	AgreementNotSpecific (d)	Defendant end of D11
F6	SecurityMeasures (p)	Any point of D13 except extreme Defendant end
F7	BroughtTools (p)	D8
F8	CompetitiveAdvantage (p)	Plaintiff end of D1
F10	SecretsDisclosedOutsiders (d)	Any point of D3 except extreme Plaintiff end
F11	VerticalKnowledge (d)	Defendant end of D2
F12	OutsiderDisclosuresRestricted (p)	Plaintiff end of D4
F13	NoncompetitionAgreement (p)	Plaintiff end of D7
F14	RestrictedMaterialsUsed (p)	Not in HYPO
F15	UniqueProduct (p)	Not in HYPO
F16	InfoReverseEngineerable (d)	Not in HYPO
F17	InfoIndependentlyGenerated (d)	Not in HYPO
F18	IdenticalProducts (p)	Not in HYPO
F19	NoSecurityMeasures (d)	Defendant end of D13
F20	InfoKnownToCompetitors (d)	Not in HYPO
F21	KnewInfoConfidential (p)	Not in HYPO
F22	InvasiveTechniques (p)	Not in HYPO
F23	WaiverOfConfidentiality (d)	Not in HYPO
F24	InfoObtainableElsewhere (d)	Not in HYPO
F25	InfoReverseEngineered (d)	Not in HYPO
F26	Deception (p)	Not in HYPO
F27	DisclosureInPublicForum (d)	Point on D3 towards Defendant end

Table 1 Base Level Factors in CATO

extreme pro-plaintiff point. F15 and F18 are harder to locate on this dimension, and so a second new dimension may be required for them. The important thing is that there is a mapping from dimension points to factors, which would allow the HYPO cases to be expressed as factors (the missing dimension D5 is used only to invalidate agreements). In general it is possible to see factors as points (or ranges) on HYPO dimensions. But this does mean that the *analyst* decides at what point the dimension ceases to favour the plaintiff and begins to favour the defendant, rather than this question forming part of the argumentation.

Representing cases as bundles of factors make it very easy to discover differences between cases. As noted in [144] this means that when comparing cases the factors can be placed in one of seven partitions:

- A Plaintiff factors in both current case and precedent.
- B Defendant factors in both current case and precedent.
- C Plaintiff factors in current case not in precedent.
- D Defendant factors in current not in precedent case.
- E Defendant factors in precedent case not in current case.
- F Plaintiff factors in precedent not in current case.
- G Factors (both plaintiff and defendant) not in either current case or precedent

This partitioning was first used in [15] from which Figure 3 is taken.

(A) and (B) represent what is similar in the cases. For the precedent case to serve as a precedent for current case there must be at least one factor in at least one of these partitions. (C) and (D) represent aspects in which current case is

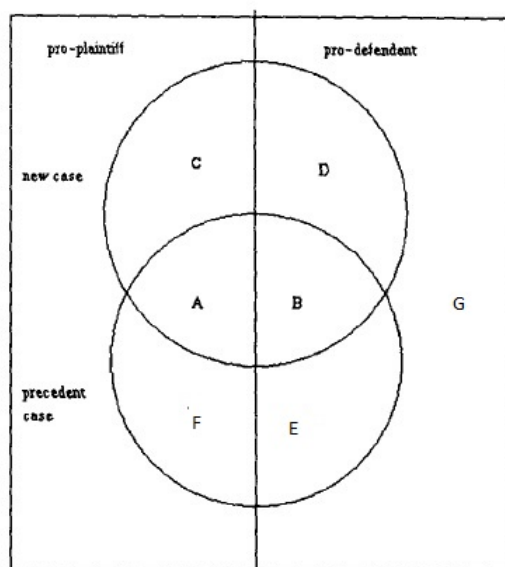


Fig. 3 Partitioning CATO factors from [15]

stronger for the plaintiff than the precedent, and so cannot be used to distinguish the precedent. (E) and (F) represent aspects in which current case is weaker for the plaintiff than the precedent, and so can be used to distinguish the precedent. (G) contains the factors which are not relevant in the comparison. Thus the law students will be advised to consider only (E) and (F) when looking for distinguishing factors that the defendant can use.

4.1.2 Downplaying and Emphasising Distinctions

Having identified which factors could serve as a distinction, the question arises as to which will be the basis of good distinctions. The answer is provided by CATO's second innovation: the introduction of abstract factors, and the organisation of factors into a factor hierarchy. The factor hierarchies of CATO are shown in Figure 4. Note that the plural is used because the root of a factor hierarchy is a legal issue, and CATO recognises five issues giving rise to five hierarchies. The hierarchies may have factors in common, but the issues are kept separate. It is important to recognise that these hierarchies are not is-a hierarchies. Like base level factors, abstract factors are also either present or absent in a case, and always favour the same side. Their children are *reasons* for their presence (if they favour the same side), or absence (if they favour different sides). Where there are reasons for and against the conflict must be resolved: sometimes the resolution is obvious: for example a waiver of confidentiality (F26) will cancel the factor F4 (Agreed not to disclose) so that the abstract factor F121 (Express confidentiality agreement) is not present. In other cases, such as F111 (Questionable Means), if both F26 (Deception) and F25 (Information Reverse Engineered) are present, it may be

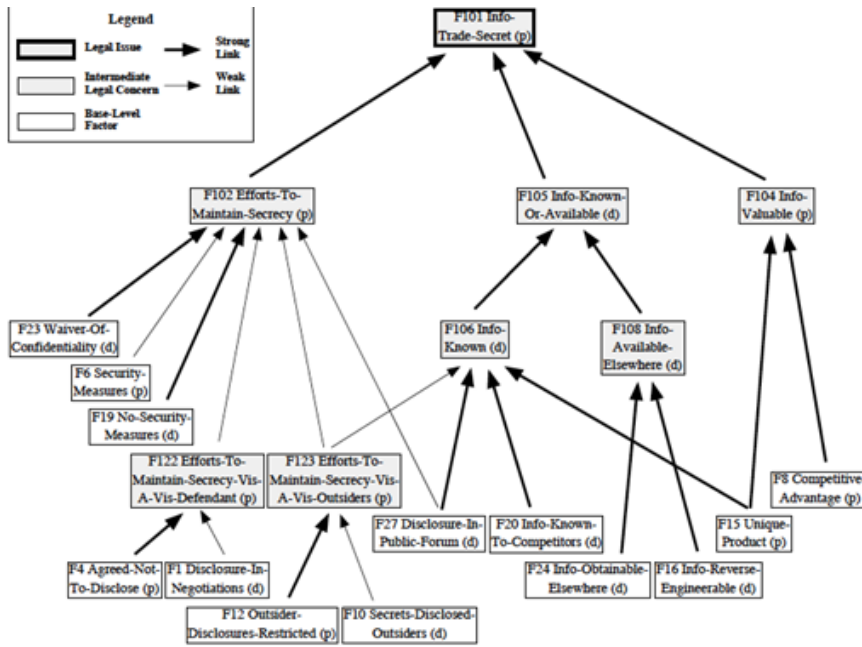


Figure 3-2: CATO's Factor Hierarchy, part 1

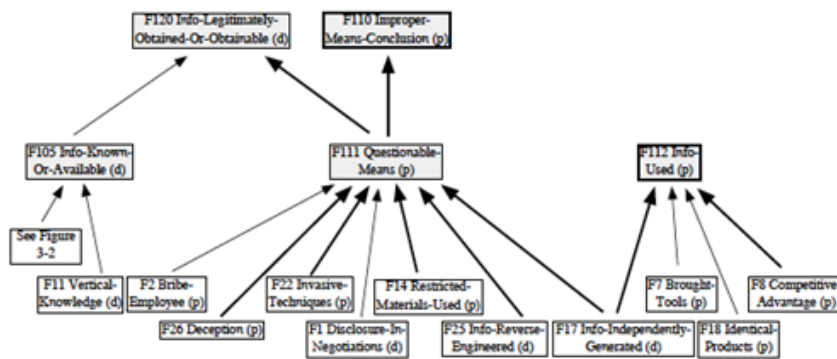
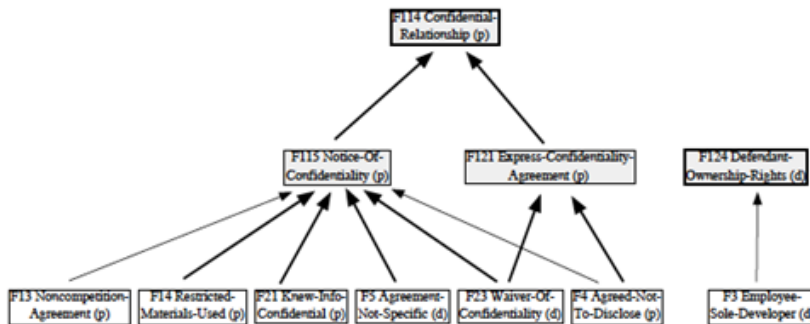


Figure 3-3: CATO's Factor Hierarchy, part 2

Fig. 4 CATO Abstract Factor Hierarchy from [10]

difficult to determine whether the abstract factor is present or not. Indeed there may be no general answer, and the particular case facts may need to be considered.

Thus a consideration of the factor hierarchy can tell us the significance of a distinction. If the “missing” factor has a sibling present in the case this can be used as an alternative way of establishing the presence of the abstract factor which represents a strength (when they favour the same side) or, when they favour different sides, as the reason for the absence of an abstract factor which would otherwise represent a weakness. This is how we can downplay a distinction. On the other hand if any additional factors are no better than cousins, or even pertain to an entirely different issue, the distinction is significant and can be emphasised. These different ways of making and challenging distinctions were expressed as argumentation schemes in [144].

4.1.3 *Argument Moves*

A third influential feature of CATO was its identification of a set of *argument moves*, several based on various operations carried out in HYPO. Although other sets of argument moves have been proposed, such as those used in CABARET [134], the CATO set has proved enduringly useful, especially when attempting to identify a procedure, such as [33] or when devising a dialogue game to undertake the reasoning [143]. The argument moves used in CATO (as listed in [10], Figure 2.2) are:

- Analogising a case to a past case with a favourable outcome;
- Distinguishing a case with an unfavourable outcome;
- Downplaying the significance of a decision;
- Emphasising the significance of a distinction;
- Citing a favourable case to emphasise strengths;
- Citing a favourable case to argue that weaknesses are not fatal;
- Citing a more on point counterexample to a case cited by an opponent;
- Citing an as on point counter example to a case cited by an opponent.

Not all of these moves are new in CATO: some can be found in HYPO. But the third and the fourth are new - made available by the factor hierarchy - and it is this list that is usually cited in the later work mentioned above. The moves are also important in identifying the critical questions [141] posed when using the argumentation schemes in [144].

4.1.4 *CATO summary*

CATO has been extremely influential on subsequent developments in AI and Law. Even when HYPO is cited as the key influence, the actual approach taken often owes as much to CATO as to HYPO. All three aspects introduced in CATO have had influence:

- CATO style factors, and the representation of cases as bundles of binary factors have been the predominant form of representation of cases in work such as dialogue [108], theory construction [53] and argumentation [144] and are the starting point for formalisations [91].

- The factor hierarchy and its introduction of intermediate predicates have been found useful in works such as [63], [25], [95] and [78]. The idea has also been used in work on methodologies for representing cases such as [7].
- Argument moves, in particular the set identified by CATO, have also influenced AI and Law work on dialogues for reasoning about cases [143].

Also heavily influenced by CATO, was the work of Ashley with another student, Stefanie Brüninghaus, which continued the exploration of reasoning with Trade Secrets law cases. This will be discussed in the next section.

4.2 IBP

Brüninghaus worked with Ashley on two projects, SMILE (*SMart Index Learner*) and IBP (*Issue-Based Prediction*) both of which are described in a paper included in this special issue [21]. The two programs were supposed to act at different ends of the case based reasoning process: SMILE would take cases input as natural language and identify which factors were present, and then IBP would deploy CATO style reasoning, adapted for prediction rather than teaching, to predict the outcome of the case. Empirical evaluations are also reported in [21], showing a strong performance from IBP, but a rather weaker performance from SMILE. Here we will concentrate on IBP, since SMILE draws mainly on classification techniques rather than the case based reasoning techniques we are focusing on here, albeit that SMILE was attempting to ascribe CATO's factors to cases.

4.2.1 Logical Model

The key modification necessary for converting CATO into a program capable of predicting case outcomes is to combine the five issue hierarchies into a single structure. As noted in section 4.1.2, and as is evident from Figure 4, CATO uses five distinct hierarchies. IBP organises these using what it terms a logical model, taken from the *Restatement of Torts*. This logical model is shown in Figure 5. The logical model serves the same purpose as the top layers of rules in CABARET [134]. Essentially the idea is that the overall question of whether a Trade Secret was misappropriated requires both that the information was a trade secret and that it was misappropriated. To be considered a secret, it needs to be shown both that the information was valuable and that the owner had made efforts to maintain its secrecy. Misappropriation could either be shown by the use of improper means to obtain the information, or by the breach of a confidential relationship in using the information. Thus the model has five leaf nodes, corresponding to the five hierarchies in CATO.

The 26 base level factors of CATO are thus split over the five issues, as indicated in Table 2. Notice that several factors relate to two issues, and one, F3, is not used in IBP.

IBP will now use the factors to decide which party is favoured on each of the issues, and then put these into the logical model to predict the outcome. The precise algorithm is given in Figure 2 of [21] in this issue, and is fully described there, so no details are required in this introduction. One additional feature introduced by IBP is that of *strength* of factors. *Knock Out Factors* are those “representing

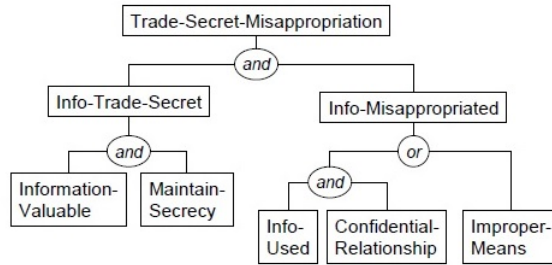


Fig. 5 IBP Logical Model from [62]

ID	Factor	CATO Issues
F1	DisclosureInNegotiations (d)	Confidential Relationship
F2	BribeEmployee (p)	Improper Means
F3	EmployeeSoleDeveloper (d)	
F4	AgreedNotToDisclose (p)	Maintain Secrecy, Confidential Relationship
F5	AgreementNotSpecific (d)	Defendant end of D11
F6	SecurityMeasures (p)	Maintain Secrecy
F7	BroughtTools (p)	Information Used, Improper Means
F8	CompetitiveAdvantage (p)	Information Valuable, Information Used
F10	SecretsDisclosedOutsiders (d)	Maintain Secrecy
F11	VerticalKnowledge (d)	Defendant end of D2
F12	OutsiderDisclosuresRestricted (p)	Maintain Secrecy
F13	NoncompetitionAgreement (p)	Confidential Relationship
F14	RestrictedMaterialsUsed (p)	Information Used, Improper Means
F15	UniqueProduct (p)	Information Valuable
F16	InfoReverseEngineerable (d)	Information Valuable
F17	InfoIndependentlyGenerated (d)	Information Used, Improper Means
F18	IdenticalProducts (p)	Information Used
F19	NoSecurityMeasures (d)	Maintain Secrecy
F20	InfoKnownToCompetitors (d)	Information Valuable
F21	KnewInfoConfidential (p)	Confidential Relationship
F22	InvasiveTechniques (p)	Improper Means
F23	WaiverOfConfidentiality (d)	Confidential Relationship
F24	InfoObtainableElsewhere (d)	Information Valuable
F25	InfoReverseEngineered (d)	Information Used, Improper Means
F26	Deception (p)	Improper Means
F27	DisclosureInPublicForum (d)	Information Valuable

Table 2 Mapping of Base Level Factors to Issues in IBP

behavior paradigmatically proscribed or encouraged under trade secret law and for which the probability that a side wins when the Factor applies is at least 80% greater than the baseline probability of the sides winning.” Conversely *Weak Factors* are those “for which the probability of the favored sides winning, given that one knows the Factor applies, is less than 20% over the baseline probability of the sides winning.” Weak factors appear to act only as supports for other factors, and IBP will not consider an issue raised if only weak factors are present. Given that factors can be related to points on HYPO style dimensions, it is not surprising that differences in strength need to be considered if we are resolving issues: the use of Knock Out and Weak factors allows this while maintaining much of the attractive

simplicity of factors as compared with dimensions. The need to consider different strengths of factors remains a current concern [6].

4.3 Empirical Results for IBP

Another feature of IBP was that it was given a thorough empirical evaluation based on 184 cases (the 148 used in CATO together with another 36 analysed specifically for IBP). This evaluation was undertaken using versions of HYPO and CATO, a version of IBP with just the logical model, and several different machine learning approaches in addition to IBP itself. IBP performed best with a 91.8% success rate, and only a single abstention. The next best performer was a Rule Learning program which achieved 88%¹⁰. The version of CATO achieved 77% (with 22 abstentions), and HYPO 67.9% with 50 abstentions. This is, of course, a little unfair on CATO and HYPO, since they were not designed to predict outcomes, and the abstentions count heavily against them. In fact, HYPO gets 93.3% of the cases for which it offers an opinion correct.

Table 3 Results, including some from [62], [65] and [7]

	correct	error	abstain	accuracy	no-abst
ANGELIC Secrets II	31	1	0	96.8	96.8
ANGELIC Refined	31	1	0	96.8	96.8
AGATHA Brute Force	30	2	0	93.7	93.7
IBP	170	15	1	91.4	91.9
AGATHA A*	29	3	0	90.6	90.6
ANGELIC Secrets	29	3	0	90.6	90.6
CATO-coin	163	30	0	89.0	89.0
Naive Bayes	161	25	0	86.5	86.5
HYPO-coin	152	34	0	81.7	81.7
ANGELIC CATO	25	7	0	78.1	78.1
CATO	152	19	22	77.8	88.8
HYPO	127	9	50	68.3	93.4
IBP-model	99	15	38	72.6	86.8

These results have been used as a basis for comparison when evaluating later programs, including AGATHA [65], a program based on the approach of [66] in this issue, and ANGELIC [3], a project which produced a methodology for encapsulating case based domains, which was applied to several domains, including Trade Secrets. These programs were forced to use only a subset of the cases, since the majority of the case analyses have unfortunately not been made available to researchers, but evaluations showed them able to achieve 90+% on these cases. The results from [3] are shown in Table 3. An earlier program based on neural networks [32] was able to reach 98+%, on a different set of cases¹¹. The full set of results for the IBP project are given as Table 3 of [21]. These results provide a

¹⁰ This program was not part of the evaluation in [3], in which Naive Bayes with 86.5% was the second best performer.

¹¹ This dataset was artificially generated, allowing for the use of a very large number (1500) cases. The dataset has been subsequently reused in [92], [102], [143] and elsewhere.

useful benchmark for other programs attempting the task, and a success rate of 90% would seem a sensible aspiration¹².

5 Integration with Rule Based Reasoning

In the 1980s and early 1990s Case Based and Rule Based systems were sometimes seen as competing approaches (although others, including both Gardner and Rissland, recognised that statutory interpretation (at least in the common law traditions) would require both rules and cases). Discussion of the two approaches was held at two of the first three ICAIL conferences, in 1987, and 1991, which was chaired by Don Berman [54]. Berman’s view seemed to be that the two approaches had different aspirations: rule based systems were more likely to produce practical applications, with rules derived either from statutes as in [132] or from an expert as in [135] in which J.C. Smith expressed his understanding as a set of rules. Berman, however, saw case based reasoning as more capable of capturing legal reasoning: “A major goal of pure AI is to represent accurately human intelligence and, therefore, to represent legal thought CBR research in the legal domain must continue” [54]. In general Rule-based systems were associated with Europeans, because of the widespread use of Prolog and logical representations of legislation, inspired by the success of [132] and the prevalence of Civil Law in continental Europe, while Case-based systems were seen as the approach of choice of US researchers, with their preference for LISP and the presence of a highly adversarial common law tradition. In the late eighties the two were seen (perhaps especially in Europe) as separate approaches and research tended to be conducted by different groups. However, exposure to each others approaches through ICAIL and elsewhere led to a greater interest in and understanding of the position of the “other side”. Increasingly in the 1990s practitioners of the different approaches were now recommending integration, with both approaches required to produce a complete system. Thus in CABARET [134] we can see rules used to provide a top level structure in which CBR can be deployed, with an agenda mechanism to control the processing of two co-equal CBR and RBR reasoners. CABARET used observations and control rules to post and order tasks on the agenda: control in CABARET is best discussed in [122]. Within the rule-based approaches, a requirement for cases to provide sufficient conditions to enable interpretation of the

¹² Although just 79% was considered enough for the work reported in [9] to make a stir in the national press and radio in the UK in October 2016: e.g *The Guardian* at www.theguardian.com/technology/2016/oct/24/artificial-intelligence-judge-university-college-london-computer-scientists and the BBC, www.bbc.co.uk/news/technology-37727387. This perhaps a reflection of the upsurge in interest in applications of AI and Law in the UK. See for example the 2016 Law Society of England and Wales panel at www.youtube.com/watch?v=8jPB-4Y3jLg and a special edition of the BBC Radio 4 programme *Law in Action* at www.bbc.co.uk/programmes/b07dlxmj. The program in [9] took the text of published judgements from the European Court of Human Rights (i.e material produced *ex post*) acting as proxies for the unavailable applications and briefs. As the authors observe, *ex ante* documents would be the input to an actual prediction system. The system produced a binary classification into *violation* and *no violation*. The task in [9] is thus more like that of SMILE+IBP (although SMILE was looking for individual factors rather than the outcome) which, because of the difficulty of the SMILE task compared with the IBP task, performs much closer to 70%. SMILE, moreover, worked on *squibs*, summaries of the cases, and although these were taken from, and used the language of, the decision, did not contain phrases related to the outcome of the case. All websites accessed March 2017.

terms of legislation was recognised in the work of proponents of rule based systems such as [31]¹³. A very real advance in allowing the approaches to be integrated was, however, made by [108] which provided an elegant way of moving from cases to a set of rules. This paper will be discussed in the next section.

5.1 Central Idea of Praken and Sartor (1998)

A number of different issues are addressed in [108], but here I will focus on its most influential idea, the movement from precedents represented as bundles of factors to sets of rules. The main aim of the translation was to enable the precedent to be used in a formal dialogue game, but the rules can equally be used as the basis of an executable logic program. Once rewritten as rules, the set of precedents are available in a form convenient for use in argumentation generally, as knowledge bases, and as the basis of formalisations of the reasoning involved.

The starting point is a set of cases represented as sets of factors together with their outcome. The factors can be divided into pro-plaintiff factors and pro-defendant factors, with each factor representing a reason to decide the case for the side they favour. It is assumed that an additional factor for a party will always strengthen the reasons to decide for that party. This may not be true in general (see [105]), but is true of CATO, and can be imposed as a constraint on what is acceptable as a factor. Now the strongest pro-plaintiff reason will be the set of all the pro-plaintiff factors present in the case (F_p) and the strongest pro-defendant reason is the set of all the pro-defendant factors present in the case (F_d). We can express this as two rules:

r1: $F_p \rightarrow p$
 r2: $F_d \rightarrow d$

where p and d represent decisions for the plaintiff and defendant respectively. Now the outcome of the case will indicate which of these reasons was preferred, and so we add a rule expressing a priority between these two rules, so that

r3: $r1 \succ r2$

represents a preference for F_p over F_d , and holds if the plaintiff won the case. Each of the precedent cases can be represent as three rules in this way, so that the whole case base can be rewritten as a set of rules to find for the plaintiff, a set of rules to find for the defendant, and an incomplete set of preferences between them. The result is a knowledge base with many similarities to the Reason Based Logic of Hage [83] and [84], further developed with his student Bart Verheij [138] and [140], but the timing and simplicity of the technique of [108], together with its strong connection to case based reasoning systems, meant that it made a more lasting impression on the AI and Law community.

The major limitation with the resulting knowledge base is that many possible rule conflicts will not be resolved by the priorities derived from the precedents. As noted in [4], the 26 factors of CATO, split evenly as they are between plaintiff and defendant, give rise to more than 8000 sets of factors favouring each of the sides,

¹³ In many ways this is just a return to the very early approach of [72] which began with rules, but which turned to cases when “the rules ran out”.

and in excess of 67 million potential comparisons. Even allowing for subsumption of proper subsets, the number of required priorities to ensure that a conflict can be resolved with precedential authority will exceed the cases available (and most probably all the cases that have ever been decided). There is, however, a desire to go beyond what is contained in the set of precedent cases, so that we can say something about cases which are not decidable *a fortiori* with respect to the precedents. In [108] this is achieved through *rule broadening*, allowing one or more factors to be removed from the antecedent, but then maintaining the priority of the original rule as established in the case from which it derived. The problem is that this lacks justification, unless or until endorsed by the judges in an actual case. Essentially, if the precedent does not provide a definitive answer, it can be distinguished by pointing to a factor missing from the plaintiff rule, or an additional factor in the rule in the precedent, and we need to rely on the user to say whether the distinction blocks the argument or not. Broadening, and several kinds of distinction, are fully discussed in [108]. Although this is a difficult problem, addressed differently in [91], discussed in the following subsection, and in the value based theory construction approaches described in section 6, the techniques of [108] at least enable a clean expression of the problem and potential solutions to it.

5.2 Formalisation of Precedential Reasoning

Although [108] provides a formal dialogue game which can be regarded as a formalization of the aspects of legal theories on judicial reasoning discussed in that paper, it is very much tied to the dialogical context, and so to a process of argumentation. Horty's main aim was to answer the question *How is it, exactly, that precedents constrain future decisions?*, and attempts to give a logic of precedent. Horty pursued this through a series of papers beginning with [87], moving through [88], [89] and [90] to its culmination in the paper included in this special issue, [91]. Horty is motivated by a desire to counter the arguments of [14]. As Horty explains in [88], Alexander (who is working in Law rather than AI and Law) had identified three different models of precedent:

- the *natural model*, by which “a precedent decision might figure into the reasoning of a court in its attempt to reach the correct decision in a current case; but on the natural model, this is the extent of precedential constraint” (p19);
- the *rule model*: by which precedents include a “rule that carries the precedential constraint. Constraint by precedent is just constraint by rules; a constrained court must apply the rules of precedent cases in reaching current decisions ... There is no room for narrowing the rule, or distinguishing the current case from the precedent;” (p20);
- the *result model* by which “a precedent controls all and only *a fortiori* cases - that is, all and only those cases that are as least as strong for the winning side of the precedent as the precedent case itself.” (p20).

In [88] Horty wishes to defend the result model against the rule model advocated by Alexander. Both Horty and Alexander reject the natural model. To defend the result model, Horty draws extensively on AI and Law research, citing amongst others, [18] and [19], [10] and [108].

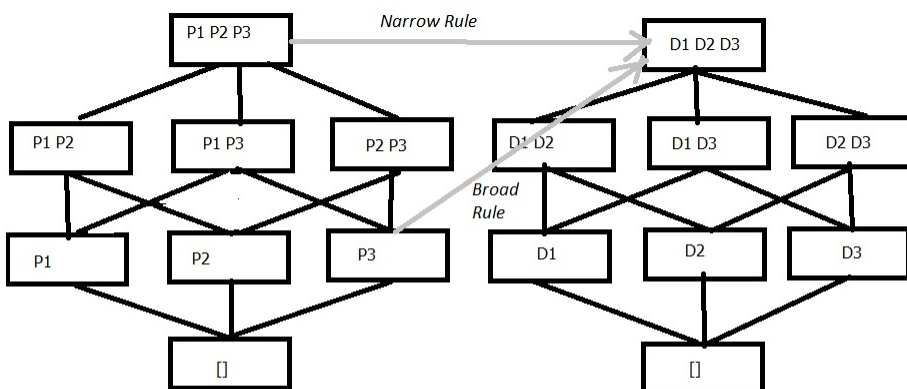


Fig. 6 Factor Lattice for domain with 3 pro-P and 3 pro-D factors. Given a precedent with all six factors present decided for the plaintiff, the plaintiff rule can be given the narrow interpretation of [108], or several broad interpretations as in [91], one of which is shown.

The main idea of Horty is to represent precedents in rules, in much the same way as [108], but with one very important difference. Suppose the case was decided for the plaintiff. Now the defendant's reason, in both [108] and [91] will be the strongest possible, that is the set of all pro-defendant factors present in the case, but the plaintiff's reason is treated differently. Whereas in [108] it is the strongest pro-plaintiff reason, that is all the pro-plaintiff factors present in the case, Horty reasons that a weaker reason may have been enough to defeat the defendant. Thus the plaintiff rule in [91] can be a subset of the factors present in the case favouring the *winner*. This makes the resulting rule stronger, because it constrains future cases to a greater extent. The effect is the same as rule broadening, but it is done at the theory level rather than the case level, and is supposed to apply to all the cases in the case base. The situation is illustrated using the pictorial notation of [34] (which followed [108] in using narrow preferences) in Figure 6.

The figure shows a Factor Lattice for a domain with six factors: three pro-plaintiff and three pro-defendant. These are grouped into two lattices one pro-plaintiff and one pro-defendant. These contain every possible combination of factors favouring the appropriate party, arranged in a partial order embodying the assumption that more factors beat fewer factors. A precedent provides a link between one node from each lattice, and so enables comparison between sets of factors from the two different lattices. Thus a narrow interpretation of the precedent with all six factors decided for the plaintiff, (shown in Figure 6) will indicate a preference for $\{P1, P2, P3\}$ over $\{D1, D2, D3\}$. This will not constrain any of the subsets. There are, however, several broad interpretations possible in the manner of [91]: any of the two member or single member sets could be said to be preferred to $\{D1, D2, D3\}$, one of which, $\{P3\}$, is shown in Figure 6. This means that $\{P3\}$ is

preferred to every combination of pro-defendant factors (since {D1 D2 D3} is preferred to every subset). Now only three pro-plaintiff subsets ({P1,P2}, {P1} and {P2}) are not constrained, and so await future decisions. For example a case with factors {P1, P2, D1} would, if decided for the defendant, now fully determine all subsequent cases. This would incidentally mean that {P3} was preferred to {P1, P2}, but this is not significant, since these subsets are never competing.

With these mechanisms Horty can formally characterise what it is for a case base with the given (broad or less broad) interpretations to be consistent. Now cases must be decided so as to maintain consistency of the case base: this may constrain the outcome, or just how broad the rule taken for the winner of the case can be. The open question is how the subset of pro-winner factors is to be determined. Horty does not specify this, although one suggestion might be that it should be the *ratio decedendi* of the case (see [60]), which can often be discovered by an examination of the text of the decision. Horty also envisages the subset being determined when the precedent is decided. A possible alternative might be that the theory is built anew for each new case, the constraint being that a consistent theory explaining the previous precedents would need to be produced. This might well, especially when there are relatively few precedents available, lead to some considerable revision of the rules to be taken from the precedent cases, but as the number of precedents increases the scope to revise rules while maintaining consistency with cases already decided will decrease. This view of reasoning with cases has some strong similarities with the theory construction approaches found in [98] and [53], and with the notion of a life cycle of case law found in [94] in which a period of turbulence is followed by a period of stability.

Horty is the starting point for [110], who likes the rule based representation of cases, especially because it is able to maintain the venerable distinction between *ratio decedendi* and *obiter dicta*. Rigoni associates Horty's broadened rules with *ratio decedendi*. Rigoni, however, offers a number of improvements to Horty's account. First he allows for multiple rules to be associated with a single precedent in order to handle *over-determined cases*. This would also allow the incorporation of *obiter dicta*. Second he recognises that not all precedents serve the same purpose: identifying a class of precedents he calls *framework cases*. His example of a framework precedent is *Lemon v. Kurtzman*¹⁴. Rigoni summarises the case:

In that case the US Supreme Court addressed the question of whether Pennsylvania's and Rhode Island's statutes that provided money to religious primary schools subject to state oversight violated the Establishment Clause of the First Amendment. The court introduced a three-pronged test and ultimately ruled that both programs did violate the Establishment Clause.

Rigoni's insight is that not all precedents express preferences: some rather supply tests which provide a framework for deciding the cases, in a manner similar to the use of the statute in CABARET and the Restatement of Torts in CATO. The three pronged test can thus be seen as three issues to be considered in *Lemon's* domain. Rigoni provides a way of accommodating framework cases in a Horty-style formalisation. For Rigoni's approach, the reader is referred to Rigoni's paper in this special issue.

¹⁴ *Lemon v. Kurtzman* 403 U.S. 602 (1971).

It is a feature of HYPO and CATO that all the precedent cases play the same role: to express a preference or, in HYPO, to constrain the n-dimensional space in some way. Thus cases are analysed against a framework of dimensions and factors developed with respect to the whole set of cases. In practice, however, the dimensions, factors and issues have developed over time, and some of the cases that introduced a factor, issue or dimension may well be in the case base. For example, if we consider the cases related to the Automobile Exception to the Fourth Amendment described in [115], [38] and [7], we see that no mention was made of privacy issues in the case which is seen as introducing the exception, *Carroll v US*¹⁵. The need to consider privacy is stressed in *South Dakota v. Opperman*¹⁶. Thus the issue of privacy must have been introduced in *Opperman*, or some earlier framework precedent. The much discussed *California v Carney*¹⁷, which, especially in the oral argument ([115], [22] and [2]) was important for its identification of the factors needed to distinguish between an automobile being used as a vehicle and an automobile being used as a residence. This was extensively discussed in the Oral Argument for that case ([115], [2]) and a number of new factors made their way into the decision. Such factors are rarely considered as present in previous cases, and so it may be a mistake to analyse earlier cases in terms of dimensions, factors and issues introduced subsequently. Understanding a body of cases as a *sequence*, although the topic of [86] and [128], is perhaps a topic which merits further exploration. Keeping a representation current in the face of changes in case law is also explored in [5].

5.3 Argumentation Schemes

The motivation for expressing a set of precedents as rules in [108] was so that the precedents could be used to deploy arguments in a dialogue game. At the time, following [73], dialogue games were a popular way of expressing legal procedures, such as particular forms of legal argumentation. Since their introduction to explain fallacies in propositional logic [85] and [96], such games tended to be built on an underlying rule base. An excellent overview of such systems is given in [106]. As this century has progressed, however, the use of dialogue games has tended to make way for argumentation schemes [141] and [142]. Although argument schemes had been used before in AI and Law: arguably the three ply argumentation of HYPO is a scheme, CABARET certainly uses schemes and Toulmin's argumentation scheme [137] was widely used (e.g [97]), often to drive dialogues as in [45], explicit use of argumentation schemes moved to the forefront when Walton's understanding of them in terms of premises, conclusions and critical questions, as expressed in [141], became prevalent.

A particular argumentation scheme (originally designed for value based practical reasoning) was applied to reasoning with legal cases in [81], and argumentation schemes were used to model a particular case in [75]. The explicit use of CATO's argumentation moves was carried out in [144]. These schemes were used to model particular cases in [39], and formalised, further developed and extended to include

¹⁵ *Carroll v. United States*, 267 U.S. 132 (1925).

¹⁶ *South Dakota v. Opperman*, 428 U.S. 364 (1976).

¹⁷ *California v. Carney*, 471 U.S. 386 (1985).

dimensions in [145], [29] and [109]. The last of these, in particular, offers a set of schemes designed to model HYPO/CATO style reasoning in a formal framework, ASPIC+ [101].

6 Purposes and Values

HYPO and CATO, as mentioned above treat all their cases as homogeneous. Whether considered as bundles of facts or bundles of factors, they are abstracted from all aspects of context. There is no consideration of the level of court, or the jurisdiction. Thus the HYPO cases ranged across a number of different states. Nor was there any consideration of date or sequence. HYPO cases ranged from 1845 to 1980, although the majority dated from the nineteen sixties and seventies [19]. This should not be seen as a defect: there is no reason why an AI model has to cover every aspect, and classroom discussions often use cases in the context independent fashion¹⁸. The intention of HYPO was to model one *necessary* part of actual appellate reasoning, and so it is not fair to complain that it is not *sufficient* in the real world. This aspect of HYPO was the subject of a series of critiques by Don Berman and Carole Hafner, who noted the absence of procedural context [55], consideration of purpose [56], and temporal aspects [57]. These papers were consolidated in a paper in a special issue of this journal (10:1-3) in memory of Don Berman [82]. Even though HYPO was not intended to address these aspects, it is important to tell (some) computer scientists that HYPO is not the whole of the story, and to explore how these aspects can be represented to augment the basic HYPO approach. These critiques were a little slow in being taken up, but meanwhile within the computational argumentation community interest was growing in Perelman's notion of an audience [103]. This had been the subject of [80] and was integrated with Dung's argumentation frameworks [70] in [36]. This led to a revisiting of the ideas of [56], which had argued that the preferences between factors reflected the purposes attributed to the law. These could change across times and jurisdictions [67], and so explain why similar cases might be decided differently at different times and in different places. As such the resolution of the conflicts could be seen as relative to the audiences to which the laws and arguments were addressed, and could be made computational using the formal characterisation of audiences in [36], that is as an ordering of the social values to which the arguments related, by equating the notion of purpose with the promotion of social values as in [37] and [41].

6.1 Value Based Theory Construction

Interest in [56] and examination of the purpose of laws, interpreted as the values they were intended to promote and protect, was revived by a group of three articles that appeared in *Artificial Intelligence and Law*, volume 10 numbers 1-3: [35], [104]

¹⁸ In fact the influential critique in [56] was focused on three cases: *Pierson v Post* (New York, 1805), *Keeble v Hickergill* (an English Queen's Bench case from 1707) and *Young v Hitchens* (an English Queen's Bench case of 1844). The critique is itself thus made using three cases from different jurisdictions, with very different dates, presented out of sequence. None the less these cases are often used when teaching property law in the US.

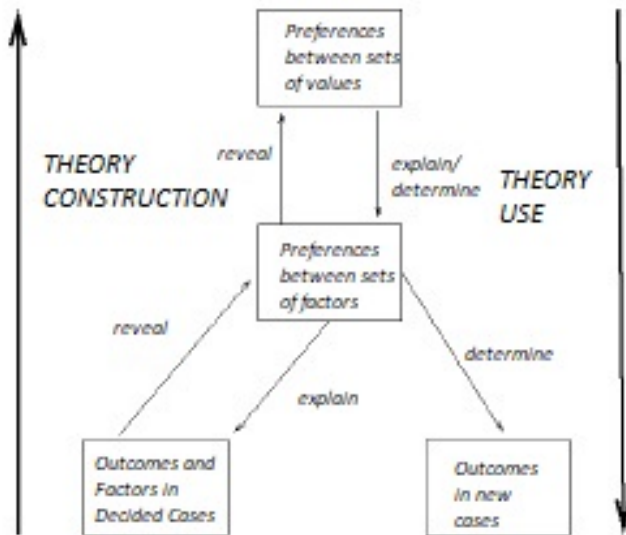


Fig. 7 Value Based Theory Construction from [52]

and [129]. Following this Sartor and Bench-Capon developed the ideas in a series of papers, including [51] and [52], culminating in [53]. These papers gave an account of theory construction with disagreements resolved according to the value preferences of the audience (here society, as represented by the judges).

The overall idea is shown in Figure 7. The starting point is at the bottom left where we have a set of cases (described as bundles of factors) and their outcomes. These outcomes reveal preferences between sets of factors, which will explain the case outcomes. The preferences are essentially the priority rules of [108]. Each factor is associated with a value: a value that will be promoted by deciding the case for the side favoured by that factor. This enables us to rewrite sets of factors as sets of values, and transfer the preferences to these sets of values. The heart of the theory is this set of value preferences. This takes us to the top of the diagram, and we have constructed a theory and so we can start to apply the theory and work our way down the right hand side. The value preferences determine and explain preferences between sets of factors which are not to be found in the existing precedents. These preferences can then be used to determine the outcomes of as yet undecided cases which have these sets of factors.

The theories are constructed using a set of operators, defined in [53]. The theory starts from nothing and construction begins by including a case from the background. This will bring with it a set of factors and their associated values. Each factor is associated with a simple rule stating that the factor is a reason to decide for the side favoured by the factor. The theory can then be extended by including more cases, combining simple rules into rules with multiple antecedents and establishing preferences between rules and between values. This continues until the theory can be applied to give an outcome for the case under consideration. At this point the onus moves to the other party who must attempt to extend

the theory to produce a better theory with an outcome for its favoured side, whereupon, it is again the turn of the original side. This process of extending and refining the theory continues until there is no possible extension of the theory which changes the outcome. The paper illustrates the process with the wild animals cases introduced in [56]. The paper also includes a discussion of how the theory operators relate to the argument moves of CATO.

This approach was tested empirically in [66] (included in this issue) and [65]. The paper included in this special issue explored the use of CATE (CAse Theory Editor) in a series of experiments intended to explore a number of issues relating to the theories constructed using the operators of [53], including how the theories should be constructed, how sets of values should be compared, and the representation of cases using structured values (which are akin to dimensions) as opposed to factors. In CATE, the construction of theories is done by the user, supported by the CATE toolset. The second paper described AGATHA (Argument Agent for Theory Automation) which is designed to automate the theory construction process, by constructing the theory first as a search over the space of possible theories, and then as a two player dialogue game (which could be played with the AGATHA program playing both sides). A set of search operators and argument moves are defined in terms of the theory constructors and the resulting theories are evaluated according to their explanatory power and their simplicity. The search or game continues until it is not possible to produce a better theory. Several search methods were investigated: brute force and heuristic search using A* and adversarial search using α/β pruning. The results proved to be good: they were reported in [65]:

AGATHA produces better theories than the hand constructed theories reported in [66], and theories comparable in explanatory power to the best performing reported technique, IBP [21]. Note also that AGATHA can be used even when there is no accepted structural model of the domain, whereas IBP relies on using the structure provided by the Restatement of Torts.

Table 3 of this introduction includes the results of some versions of AGATHA.

6.2 Other Uses of Values

Apart from their pivotal role in the approach to theory construction advocated in [51], values have been applied to case based reasoning in several other ways. Their original use, to tailor an argumentation framework to the preferences of an audience using Value based Argumentation Frameworks (VAF) [36] was applied to reasoning with cases in [25], [144] and [26], but recent work by Atkinson and her colleagues has placed moved away from this use of values in [4] and [8]. They continue to use the VAF approach for legislative law making and e-democracy [28], [27] and [42], and in the justification of norms [47].

Values have also been used for reasoning with cases in [16], which recognised that values could be used to justify rules (their traditional role in VAF based systems) and also the inclusion of particular antecedents in rules, and in [78], from which developed a different formalism intended to capture dynamic aspects of reasoning with legal cases [79]. The most recent work in this strand is [77].

Finally, another approach to using values in the representation of cases can be found in [139].

One issue relating to values that has yet to be fully explored is that it seems that they can be used in a variety of ways. Sometimes they appear to give rise to conflicting positions which can be resolved by preferring one value to another, which is essentially the approach of VAF based systems. At others they appear to be better regarded a *sets* of reasons, for which issues such as accrual are pertinent ([105], [100] and [49]), which seems applicable in domains such as US Trade Secrets, when a number of different aspects of the case must be considered. Also, however, there seem to be situations where a balance ([93] and [76]) must be struck between two values: it cannot be said that one is preferred to another, but rather both must be respected to an appropriate degree. This is perhaps the case with the Automobile Exception cases, where a balance must be struck between exigency and privacy, [48]. Unpicking these different aspects is something that AI and Law may need to address in future work.

7 Back to Dimensions

The work we have been discussing has tended to adopt CATO's version of factors, features either present or absent and always favouring the same side, rather than HYPO's dimensions, which represent a range of values, and may favour either side according to where on the dimension the case lies. This is understandable and the use of such factors is a very sensible simplification for the purpose which CATO was intended to address, the teaching of distinguishing cases. Although there are other sensible simplifications (see, for example, [124]), CATO's notion of factors has been widely adopted and this has enabled understanding of reasoning with legal cases to make the significant progress represented by the papers included in this special issue. It remains the case, however, that factors do represent a simplification and so, if we are to make further progress, it will be necessary to improve our understanding of dimensions. The differences between dimensions and factors are discussed in [117] in this special issue: in the next section I will look at some more recent attempts to make use of dimensions.

7.1 Factors and Dimensions

In the Restatement of Torts (quoted in [25]) we find (*italics mine*):

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.

Although we can identify all of these "factors to consider" with factors in CATO, the language here is, as the italicised phrases show, not really consistent

with the all or nothing, present or absent, nature of CATO-style factors. Each of them seems to require some kind of quantitative estimate of *how much?*. Thus F16, *Reverse Engineerable*, a CATO factor which has presented a number of difficulties for subsequent systems attempting to use CATO's analysis, such as that reported in [3], cries out for a more nuanced representation. *Reverse Engineerable* could mean that a person could build a version after a cursory inspection, or could require many person years of expert effort. Arguably, anything is reverse engineerable with sufficient effort, expertise and ingenuity. There are three ways in which issues requiring dimensions can arise:

1. Whether the facts favour the plaintiff or defendant: if a hunter is chasing a wild animal, how close does he have to get to count as possessing the animal? The famous case of *Pierson v Post* turned on precisely this point, as discussed in [50].
2. Whether the factor should or should not be ascribed to the case: at what level of effort does a product cease to be considered reverse engineerable?
3. In comparisons: for example, thinking in terms of values, is the weak promotion of a preferred value to be preferred to the strong promotion of a lesser value [130]? This issue is also explored empirically in [66], and more recently in [6].

The first two of these concern what can be argued about. When using factors, the decisions are made by the analyst and once made cannot be debated when attempting to resolve the case. (1) is effectively decided when defining the factors: thus whether Justinian, who demands actual physical possession for the animal to count as *caught* should be chosen over other authorities who make less stringent demands, such as mortal wounding or certain capture, is something that the analyst needs to be told along with the existence of *caught* as a factor. Once the definition has been chosen, ascription to cases is relatively easy. (2) relates to the analysis of particular cases. Given that the information was not in fact reverse engineered (in which case F25 would apply instead), the analyst must decide whether “the *ease or difficulty* with which the information could be properly acquired or duplicated by others” is such that F16 can be deemed to be present. (3) rather concerns computation, and allows mapping onto numbers to facilitate accrual and comparison. We need to be able to argue about whether a factor applies to a case and which side is favoured, but using CATO-style factors these issues are resolved by the analyst and so outside the scope of systems which begin with cases represented as bundles of factors.

7.2 Current Use of Dimension for Computation

The computation aspect is key in [66]. The main concern in that paper is values, but there is a section on what it terms *structured values*. The 26 CATO factors are associated with one of five values: *CA* for respect for confidentiality agreements, *QM* for questionable methods used by the defendant, *LM* for legitimate methods used by the defendant, *RE* for reasonable efforts to maintain secrecy of the part of the plaintiff, and *MW* for material worth of the information.. A structured value runs from an extreme pro-plaintiff point to an extreme pro-defendant point, and the various CATO factors relating to the value are assigned positions along this line according to how strongly they promote or demote the value. They are thus

Table 4 Dimensions in [6], Their Values and Their Factors.

	Agreement	Dubious	Legitimate	Measures	Worth	Disclosure	Availability
	CA	QM	LM	RE	MW	RE	LM
10	F4	F22		F6	F8		
9		F2					
8	F13	F26	F7		F7	F12	
7	F21	F14					
6	F14				F18		F15
5							
4						F1	
3	F5		F25		F16		F20
2							
1	F23		F17			F10	
0			F3	F19	F11	F27	F24

dimensions in all but name. The values can themselves be ordered, and to the contribution of the factor becomes a function of the importance of the value and the extent to which the particular factor promotes it. The empirical consequences of the additional degrees of freedom this enables are given as part of the results in [6]¹⁹ and Table 3 of this paper.

This approach is also adopted in [43] which provides dimensions for the wild animal cases of [56] and [6] which applies the same process to CATO's factors. In [6] seven dimensions are used rather than the five structured values of [66], and the dimension runs from 10 (extreme pro-plaintiff) to 0 (extreme pro-defendant) rather than from +10 to -10, as used by [66].

The seven dimensions in [6] are:

- **Agreement:** This relates to the existence of explicit confidentiality and non-disclosure agreements.
- **Dubious:** This considers whether the means used by the defendant were illegal or otherwise dubious or questionable
- **Legitimate** This covers the independent discovery of the information.
- **Measures:** This relates to the security measures used by the plaintiff.
- **Worth:** This relates to the value of the information, in terms of the time and effort that it might save, and the value it would add to the defendant's product.
- **Disclosure:** This relates to the number of people to whom the information and been disclosed and the circumstances of such disclosures.
- **Availability:** This considers how readily available the information was to the defendant.

The distribution of the 26 CATO factors is shown in Table 4.

These dimensions allowed computation of predictions of case outcomes using propositions with truth values between 0 and 1 to represent the dimension points, and interpreting the logical operators in the manner of fuzzy logic [146], as proposed in [46]. The results of the dimensional version of Angelic Secrets reported in [6] showed an improvement on previously reported versions which had used boolean factors [3]. This seemed to centre in particular on a better handling of the vague factor F16, reverse engineerability. Two other cases with factors which it had been suggested in [7] were wrongly ascribed, were wrongly decided by the

¹⁹ Several papers cited in this section report work carried out by Latifa Al-Abdulkarim as part of her PhD project [1]. This project developed the ANGELIC (ADF for kNowledGe Encapsulation of Legal Information from Cases) methodology, which was applied to several domains including US Trade Secrets (Angelic Secrets) and the wild animals cases (Angelic Animals).

program, but when the factors in these cases were corrected in the ways suggested in that earlier paper a success rate of 96.8% was achieved. This suggests that it is possible to achieve better results if we allow for degrees of support, rather than insisting on the all or nothing choices required by the use of factors. Note, however, that even with the dimensional approach of [6], the description of the cases in terms of factors/dimensions is still a matter for the analyst: it cannot be disputed within the program. This relates to issues (1) and (2) above and will be discussed in section 7.3. Note, however, that there is no obligation for the computer system to cover the whole process: as discussed in [8] a system may be scoped so as to start with factors or dimensions. It may be preferable to allow an analyst to represent the case in terms of dimensions and factors rather than attempt to enable their derivation from the brute facts. There is a degree of trade off here: see also [20] and the commentary on it by Thorne McCarty in [40].

The use of dimensions can also be used to generate case narratives, as described in [43], using methods akin to the scripts of [131]. These narratives can then be used to recast the story of the case in terms of legally significant facts, and to support improved explanations²⁰.

7.2.1 Comparison with HYPO dimensions

The seven dimensions used in [6] are somewhat different from those used in the original HYPO program as reported in [19]. This should come as no surprise: [6] was able to draw on three decades of work on representations of US Trade Secrets, above all CATO, but also the work of [21] and [66]. Table 5 shows the relation between the two sets of dimensions. In some cases we see several of HYPO's dimensions arranged on same Angelic dimension: confidentiality, for example embraces four of HYPO's (binary) dimensions. In contrast, the dimensions relating to the acquisition of the information seem underrepresented in HYPO, and the dimension Availability, corresponds to none of HYPO's dimensions, and so is omitted from Table 5. These differences are probably because the cases introduced in CATO contained several cases which related to this issue. HYPO, in contrast, contained a number of cases relating to the validity of agreements: the older cases related to validity of contracts not specifically associated with Trade Secrets. This is reflected in HYPO's dimensions, in particular include the notion of *consideration*, which relates specifically to the validity of any agreements and does not feature in either CATO or Angelic Secrets (where an agreement without consideration simply does not appear as a factor).

7.3 Dimensions as a Bridge to Facts

The transition from the facts of a case as reported using "world facts" to a set of legally significant "legal facts" has long been an issue in AI and Law, at least since [61]. The difficulty is that the world facts are diverse and peculiar to particular cases, whereas the legal facts exhibit more uniformity and have a more definite

²⁰ This also revives an idea from the past: for example [123] discussed recurring prototypical fact patterns or legal stories in terms of legal scripts. Schank himself was an invited speaker at the second ICAIL (1989).

Table 5 Dimensions in [6] and [19]

	Confidentiality Agreement	Dubious Means	OK Means	Security Measures	Material Worth	Disclosure
	CA	QM	LM	RE	MW	RE
10	Agreed Not To Disclose			Security Measures(P)	Competitive Advantage(P)	Secrets Disclosed Outsiders(P)
9		Bribe Employee			Vertical Knowledge(P)	
8	Non competition Agreement				Brought Tools	Outsider Disclosures Restricted(P)
7						
6						
5						
4						Disclosure In Negotiations
3	Agreement Not Specific					
2						
1					Vertical Knowledge(D)	Outsider Disclosures Restricted(D)
0	No Consideration		Employee Sole Developer	Security Measures(D)	Competitive Advantage(D)	Secrets Disclosed Outsiders(D)

Table 6 Summary of Statement Types. Base Level Factors, Legal Facts and Factual Conclusions all Correspond to Ranges on Dimensions

Statement Type		"Truth" Value	Justification
Verdict	V	0, 1	Entailed by II/LI
Intermediate Issue	II	0,1	Entailed by LI
Leaf 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 LI
Legal Fact	LF	0,1	$if \in FC$ meets Proof Standard
Factual Conclusion	FC	0 ...1	Entailed by IF
Intermediate Fact	IF	0 ... 1	Derived from E
Evidence	E	1	Given

relation to the issues to be decided. It is often said that once the world facts have been qualified into legal facts, the decision itself is often obvious. This clearly relates to issues (1) and (2) from section 7.1. Both these questions need to be able to be made the subject of argument.

As can be seen from the use of structured values in [66] and the Tables in this paper, especially Table 4, factors can be regarded as points (or ranges) on dimensions. So the first move is to see the legal facts of a case as a bundle of dimension points. Then a case description can be seen as a frame with slots corresponding to dimensions and fillers as points/ranges of those dimensions. This is precisely the structure advocated in [43] and [6]. It is also the method of representing cases in [109] to enable argumentation about the how the factors present in a case relate to facts and dimensions.

The matter is made more explicit in [8], which distinguishes various types of statements in the journey from evidence to verdict, with a particular emphasis on how we can move from the world facts to the legal facts. It is important not to confuse issues of fact and issues of law. Table 6 shows the different types of statement used in that paper.

Note that the points/ranges on dimensions play three roles. First, they are the conclusion of the reasoning about the world, and the input to the legal reasoning.

Thus we start with a mass of evidence from which we need to arrive at a set of well structured legal facts which will relate to the law governing the case. As such they are typically not certain but established with a certain degree of belief. The methods used here will be the standard ways in which people reason about the world: evidence takes many forms: witness testimony, documents, forensics, video footage etc. Each of these are associated with their own kind of reasoning, coherence [58], probability [136], etc, and the courts should make use of these established ways of reasoning. There is nothing distinctively *legal* in this phase, and the decisions are often made by lay juries, and often facts cannot be revisited at the Appeal stage. At this point the appropriate standard of proof ([71] and [74]) is applied. The appropriate proof standard will depend on the jurisdiction and the nature of the case: typically criminal cases impose a higher standard of proof than civil cases.

Those factual conclusions that meet the required standards become legal facts. At this point they become either true or false: all factual conclusions that meet the proof standard are “equally true”: the question is not the extent to which they are believed, but whether they are believed sufficiently to be taken as true in the legal reasoning. This is the bridge between the world and realm of law: the proof standard is the toll which gives entry to the bridge. On leaving the bridge the legal facts assume their third role: they become base level factors. As such they again become associated with a number: 1 for extreme pro-plaintiff factors, -1 for extreme pro-defendant factors and somewhere in between otherwise. Note that the number is entirely independent of the degree of belief with the factual conclusion was established: the base level factor number is determined by the dimension and the position of the factor on that dimension. Conflation of these roles can lead to problems: perhaps applying the standard of proof during legal reasoning, or allowing the degree of belief to contaminate the contribution of the factor.

All this is quite recent work and nothing like the consensus about, let alone the understanding of, reasoning with CATO-style factors has emerged with respect to dimensions. None the less, this is an active area of work, and points to a direction which may serve to enhance our understanding of reasoning with legal cases.

8 Conclusion

Reasoning with legal cases has been a central concern of AI and Law since its earliest days. As can be seen from the papers in this special issue and the many other papers referred to in this introduction, HYPO and its descendants - especially CATO - have had an enormous influence, setting the agenda for a variety of investigations carried out by different groups over three decades. As a result, the role of legal cases in moving from legal facts through intermediate concepts to issues and a final verdict is relatively well understood and has been formalised in [91] and [110]. Also it has given rise to implementations which have performed well in empirical evaluations, capable of predicting case outcomes in more than 90% of cases, as shown in [21] and [66].

All of this, however, has used CATO-style factors, and so has dealt with factors as present or absent, all or nothing, and so does not capture nuances associated with degrees of presence and degrees of support: the kind of factors suggested by the language used in the Restatement of Torts. This has led to a refocus of

HYPO's original mode of representation, using dimensions, which does allow for these notions.

There are currently a number of open questions which we would expect to be pursued in the near and middle future. These include:

- Can we modify formalisations of the sort found in [91] to accommodate factors associated with magnitudes, rather than just binary true and false?
- [91] is concerned with comparing sets of factors. What about cases where factors need to be balanced?
- Can we use dimensions to be more precise about the transition from world facts to legal facts?
- Can we exploit the structure of cases into dimensions and points/ranges on these dimensions when reasoning about evidence?
- What roles can precedents have? We have seen that there are precedents indicating preferences, precedents providing a framework of issues as in [110], and tentatively identified precedents which introduced additional factors. Are there others?
- What is the role of purposes and values? If values can play several roles, as suggested in section 6, how can these be distinguished, and how do they relate?

These are some of the questions that will need to be answered if reasoning with cases in the manner associated with HYPO and its successors is to be further developed. Thus HYPO will continue to inspire work in AI and Law for years to come.

Acknowledgments

I am very grateful to Edwina Rissland, Kevin Ashley and Katie Atkinson, for their invaluable comments on a earlier draft. Needless to say, errors that remain are all my own. I am also grateful to all the members of the AI and Law community, too numerous to mention individually, with who I have worked on, and talked about, reasoning with legal cases over the past thirty years.

References

1. L. Al-Abdulkarim. *Representation of case law for argumentative reasoning*. PhD thesis, University of Liverpool, 2016.
2. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. From oral hearing to opinion in the US Supreme Court. In *Legal Knowledge and Information Systems - JURIX 2013: The Twenty-Sixth Annual Conference*,, pages 1–10, 2013.
3. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Evaluating the use of Abstract Dialectical Frameworks to represent case law. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 156–160. ACM, 2015.
4. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Factors, issues and values: Re-visiting reasoning with cases. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 3–12. ACM, 2015.
5. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Accommodating change. *Artif. Intell. Law*, 24(4):409–427, 2016.
6. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Angelic secrets: bridging from factors to facts in US Trade Secrets. In *Legal Knowledge and Information Systems - JURIX 2016: The Twenty-Ninth Annual Conference*, pages 113–118, 2016.

7. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks. *Artificial Intelligence and Law*, 24(1):1–49, 2016.
8. L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Statement types in legal argument. In *Legal Knowledge and Information Systems - JURIX 2016: The Twenty-Ninth Annual Conference*, pages 3–12, 2016.
9. N. Aletras, D. Tsarapatsanis, D. Preotiuc-Pietro, and V. Lampos. 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, 2016.
10. V. Alevén. *Teaching case-based argumentation through a model and examples*. PhD thesis, University of Pittsburgh, 1997.
11. V. Alevén. Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment. *Artificial Intelligence*, 150(1-2):183–237, 2003.
12. V. Alevén and K. D. Ashley. What law students need to know to win. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 152–161. ACM, 1993.
13. V. Alevén and K. D. Ashley. Doing things with factors. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 31–41. ACM, 1995.
14. L. Alexander. Constrained by precedent. *Southern California Law Review*, 63:1–64, 1989.
15. M. Allen, T. Bench-Capon, and G. Staniford. A multi-agent legal argument generator. In *Proceedings of the 11th International Workshop on Database and Expert Systems Applications*, pages 1080–1084. IEEE, 2000.
16. M. Araszkievicz. Analogy, similarity and factors. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 101–105. ACM, 2011.
17. M. Araszkievicz, A. Lopatkiewicz, A. Zienkiewicz, and T. Zurek. Representation of an actual divorce dispute in the parenting plan support system. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 166–170. ACM, 2015.
18. K. D. Ashley. Toward a computational theory of arguing with precedents. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, pages 93–102. ACM, 1989.
19. K. D. Ashley. *Modeling legal arguments: Reasoning with cases and hypotheticals*. MIT press, Cambridge, Mass., 1990.
20. 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. ACM, 2009.
21. K. D. Ashley and S. Brüninghaus. Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law*, 17(2):125–165, 2009.
22. K. D. Ashley, C. Lynch, N. Pinkwart, and V. Alevén. A process model of legal argument with hypotheticals. In *Legal Knowledge and Information Systems: JURIX 2015: The Twenty-First Annual Conference*, pages 1–10, 2008.
23. K. D. Ashley and E. L. Rissland. But, see, accord: generating blue book citations in hypo. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 67–74. ACM, 1987.
24. K. Atkinson. Introduction to special issue on modelling Popov v. Hayashi. *Artificial Intelligence and Law*, 20(1):1–14, 2012.
25. K. Atkinson and T. Bench-Capon. Legal case-based reasoning as practical reasoning. *Artificial Intelligence and Law*, 13(1):93–131, 2005.
26. K. Atkinson and T. Bench-Capon. Argumentation and standards of proof. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law*, pages 107–116. ACM, 2007.
27. K. Atkinson, T. Bench-Capon, D. Cartwright, and A. Wyner. Semantic models for policy deliberation. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 81–90. ACM, 2011.
28. K. Atkinson, T. Bench-Capon, and P. McBurney. PARMENIDES: facilitating deliberation in democracies. *Artificial Intelligence and Law*, 14(4):261–275, 2006.
29. K. Atkinson, T. Bench-Capon, H. Prakken, and A. Wyner. Argumentation schemes for reasoning about factors with dimensions. In *Legal Knowledge and Information Systems - JURIX 2013: The Twenty-Sixth Annual Conference*, pages 39–48, 2013.
30. J. L. Austin and G. J. Warnock. *Sense and sensibilia*. Oxford University Press, Oxford, 1962.

31. T. Bench-Capon. Practical legal expert systems: the relation between a formalisation of law and expert knowledge. *Computers, Law and AI*, pages 191–201, 1991.
32. T. Bench-Capon. Neural networks and open texture. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 292–297. ACM, 1993.
33. T. Bench-Capon. Arguing with cases. In *Legal Knowledge and Information Systems - JURIX 1997: The Tenth Annual Conference*, pages 85–100, 1997.
34. T. Bench-Capon. Some observations on modelling case based reasoning with formal argument models. In *Proceedings of the 7th International Conference on Artificial Intelligence and Law*, pages 36–42. ACM, 1999.
35. T. Bench-Capon. The missing link revisited: The role of teleology in representing legal argument. *Artificial Intelligence and Law*, 10(1-3):79–94, 2002.
36. T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
37. T. Bench-Capon. Try to see it my way: Modelling persuasion in legal discourse. *Artificial Intelligence and Law*, 11(4):271–287, 2003.
38. T. Bench-Capon. Relating values in a series of supreme court decisions. In *Legal Knowledge and Information Systems - JURIX 2011: The Twenty-Fourth Annual Conference*, pages 13–22, 2011.
39. T. Bench-Capon. Representing Popov v Hayashi with dimensions and factors. *Artificial Intelligence and Law*, 20(1):15–35, 2012.
40. T. Bench-Capon, M. Araszkievicz, K. Ashley, K. Atkinson, F. Bex, F. Borges, D. Bourcier, P. Bourguine, J. G. Conrad, E. Francesconi, et al. A history of AI and Law in 50 papers: 25 years of the International Conference on AI and Law. *Artificial Intelligence and Law*, 20(3):215–319, 2012.
41. T. Bench-Capon, K. Atkinson, and A. Chorley. Persuasion and value in legal argument. *Journal of Logic and Computation*, 15(6):1075–1097, 2005.
42. T. Bench-Capon, K. Atkinson, and A. Wyner. Using argumentation to structure e-participation in policy making. In *Transactions on Large-Scale Data-and Knowledge-Centered Systems XVIII*, pages 1–29. Springer, 2015.
43. T. Bench-Capon and F. Bex. Cases and stories, dimensions and scripts. In *Legal Knowledge and Information Systems: JURIX 2015: The Twenty-Eighth Annual Conference*, volume 279, pages 11–20. IOS Press, 2015.
44. T. Bench-Capon, J. B. Freeman, H. Hohmann, and H. Prakken. Computational models, argumentation theories and legal practice. In C. Reed and T. Norman, editors, *Argumentation Machines*, pages 85–120. Springer, 2003.
45. T. Bench-Capon, T. Geldard, and P. Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8(2-3):233–254, 2000.
46. T. Bench-Capon and T. Gordon. Two tools for prototyping legal CBR. In *Legal Knowledge and Information Systems - JURIX 2015: The Twenty-Eighth Annual Conference*, pages 177–178, 2015.
47. T. Bench-Capon and S. Modgil. Norms and value based reasoning: justifying compliance and violation. *Artificial Intelligence and Law*, 25(1):1–36, 2017.
48. T. Bench-Capon and H. Prakken. Using argument schemes for hypothetical reasoning in law. *Artificial Intelligence and Law*, 18(2):153–174, 2010.
49. T. Bench-Capon, H. Prakken, and W. Visser. Argument schemes for two-phase democratic deliberation. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 21–30. ACM, 2011.
50. T. Bench-Capon and E. Rissland. Back to the future: Dimensions revisited. In *Legal Knowledge and Information Systems - JURIX 2001*, pages 41–52, 2001.
51. T. Bench-Capon and G. Sartor. Using values and theories to resolve disagreement in law. In *Legal Knowledge and Information Systems - JURIX 1997: The Thirteenth Annual Conference*, pages 73–84, 2000.
52. T. Bench-Capon and G. Sartor. Theory based explanation of case law domains. In *Proceedings of the 8th International Conference on Artificial Intelligence and Law*, pages 12–21. ACM, 2001.
53. T. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1):97–143, 2003.
54. D. H. Berman. Developer's choice in the legal domain: the Sisyphean journey with CBR or down hill with rules (a working paper for the case-rules panel at the third International Conference on Artificial Intelligence and Law). In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 307–309. ACM, 1991.

55. D. H. Berman and C. D. Hafner. Incorporating procedural context into a model of case-based legal reasoning. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 12–20. ACM, 1991.
56. D. H. Berman and C. D. Hafner. Representing teleological structure in case-based legal reasoning: The missing link. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pages 50–59, 1993.
57. D. H. Berman and C. D. Hafner. Understanding precedents in a temporal context of evolving legal doctrine. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 42–51. ACM, 1995.
58. F. Bex. *Arguments, stories and criminal evidence: A formal hybrid theory*, volume 92. Springer, Dordrecht, 2011.
59. L. K. Branting. Reasoning with portions of precedents. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 145–154. ACM, 1991.
60. L. K. Branting. A computational model of ratio decidendi. *Artificial Intelligence and Law*, 2(1):1–31, 1993.
61. J. Breuker and N. Den Haan. Separating world and regulation knowledge: where is the logic. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 92–97. ACM, 1991.
62. S. Bruninghaus and K. D. Ashley. Predicting outcomes of case based legal arguments. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 233–242. ACM, 2003.
63. S. Bruninghaus and K. D. Ashley. A predictive role for intermediate legal concepts. In *Legal Knowledge and Information Systems: JURIX 2003: the Sixteenth Annual Conference*, pages 153–162. IOS Press, 2003.
64. J. P. Callan, W. B. Croft, and S. M. Harding. The INQUERY retrieval system. In *Database and expert systems applications*, pages 78–83. Springer, 1992.
65. A. Chorley and T. Bench-Capon. Agatha: Using heuristic search to automate the construction of case law theories. *Artificial Intelligence and Law*, 13(1):9–51, 2005.
66. A. Chorley and T. Bench-Capon. An empirical investigation of reasoning with legal cases through theory construction and application. *Artificial Intelligence and Law*, 13(3-4):323–371, 2005.
67. G. Christie. *The notion of an ideal audience in legal argument*. Springer, Law and Philosophy Library, volume 45. Dordrecht, 2000.
68. J. G. Conrad and J. Zeleznikow. The role of evaluation in AI and Law: an examination of its different forms in the AI and Law journal. In *Proceedings of the 15th international conference on artificial intelligence and law*, pages 181–186. ACM, 2015.
69. J. J. Daniels and E. L. Rissland. Finding legally relevant passages in case opinions. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, pages 39–46. ACM, 1997.
70. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
71. 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.
72. A. Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. MIT press, Cambridge, Mass., 1987.
73. T. F. Gordon. The Pleadings Game. *Artificial Intelligence and Law*, 2(4):239–292, 1993.
74. T. F. Gordon and D. Walton. Proof burdens and standards. In *Argumentation in (A)*.
75. T. F. Gordon and D. Walton. Pierson vs. Post revisited - A reconstruction using the Carneades argumentation framework. In *Computational Models of Argument: Proceedings of COMMA 2006*, pages 208–219, 2006.
76. T. F. Gordon and D. Walton. Formalizing balancing arguments. In *Computational Models of Argument - Proceedings of COMMA 2016*, pages 327–338, 2016.
77. M. Grabmair. Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law*. ACM, In Press, 2017.
78. M. Grabmair and K. D. Ashley. Facilitating case comparison using value judgments and intermediate legal concepts. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 161–170. ACM, 2011.

79. M. Grabmair and K. D. Ashley. Using event progression to enhance purposive argumentation in the value judgment formalism. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pages 73–82. ACM, 2013.
80. F. Grasso, A. Cawsey, and R. Jones. Dialectical argumentation to solve conflicts in advice giving: a case study in the promotion of healthy nutrition. *International Journal of Human-Computer Studies*, 53(6):1077–1115, 2000.
81. K. Greenwood, T. Bench-Capon, and P. McBurney. Towards a computational account of persuasion in law. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 22–31. ACM, 2003.
82. C. D. Hafner and D. H. Berman. The role of context in case-based legal reasoning: teleological, temporal, and procedural. *Artificial Intelligence and Law*, 10(1-3):19–64, 2002.
83. J. Hage. Monological reason-based logic: a low level integration of rule-based reasoning and case-based reasoning. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 30–39. ACM, 1993.
84. J. Hage. A theory of legal reasoning and a logic to match. *Artificial Intelligence and Law*, 4(3-4):199–273, 1996.
85. C. L. Hamblin. *Fallacies*. Methuen, London, UK, 1970.
86. J. Henderson and T. Bench-Capon. Dynamic arguments in a case law domain. In *Proceedings of the 8th international conference on Artificial Intelligence and Law*, pages 60–69. ACM, 2001.
87. J. F. Horty. Precedent, deontic logic, and inheritance. In *Proceedings of the 7th International Conference on Artificial Intelligence and Law*, pages 63–72. ACM, 1999.
88. J. F. Horty. The result model of precedent. *Legal Theory*, 10(1):19–31, 2004.
89. J. F. Horty. Reasons and precedent. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 41–50. ACM, 2011.
90. J. F. Horty. Rules and reasons in the theory of precedent. *Legal Theory*, 17(1):1–33, 2011.
91. J. F. Horty and T. J. Bench-Capon. A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20(2):181–214, 2012.
92. B. Johnston and G. Governatori. Induction of defeasible logic theories in the legal domain. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 204–213. ACM, 2003.
93. M. Lauritsen. On balance. *Artificial Intelligence and Law*, 23(1):23–42, 2015.
94. E. H. Levi. An introduction to legal reasoning. *The University of Chicago Law Review*, 15(3):501–574, 1948.
95. L. Lindahl and J. Odelstad. Intermediaries and intervenients in normative systems. *Journal of Applied Logic*, 6(2):229–250, 2008.
96. J. D. Mackenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, 8(1):117–133, 1979.
97. C. C. Marshall. Representing the structure of a legal argument. In *Proceedings of the 2nd international conference on Artificial Intelligence and Law*, pages 121–127. ACM, 1989.
98. L. T. McCarty. An implementation of *Eisner v. Macomber*. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 276–286. ACM, 1995.
99. B. McLaren and K. Ashley. Case representation, acquisition, and retrieval in sirocco. In *International Conference on Case-Based Reasoning*, pages 248–262. Springer, 1999.
100. S. Modgil and T. Bench-Capon. Integrating dialectical and accrual modes of argumentation. In *Computational Models of Argument: Proceedings of COMMA 2010*, pages 335–346, 2010.
101. S. Modgil and H. Prakken. The ASPIC+ framework for structured argumentation: a tutorial. *Argument & Computation*, 5(1):31–62, 2014.
102. M. Možina, J. Žabkar, T. Bench-Capon, and I. Bratko. Argument based machine learning applied to law. *Artificial Intelligence and Law*, 13(1):53–73, 2005.
103. C. Perelman and L. Olbrechts-Tyteca. *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press, Indiana, 1971.
104. H. Prakken. An exercise in formalising teleological case-based reasoning. *Artificial Intelligence and Law*, 10(1-3):113–133, 2002.
105. H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the 10th International Conference on Artificial Intelligence and Law*, pages 85–94. ACM, 2005.

106. H. Prakken. Formal systems for persuasion dialogue. *The Knowledge Engineering Review*, 21(02):163–188, 2006.
107. H. Prakken and G. Sartor. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence and Law*, 4(3-4):331–368, 1996.
108. H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6(3-4):231–87, 1998.
109. H. Prakken, A. Wyner, T. Bench-Capon, and K. Atkinson. A formalization of argumentation schemes for legal case-based reasoning in ASPIC+. *Journal of Logic and Computation*, 25(5):1141–1166, 2015.
110. A. Rigoni. An improved factor based approach to precedential constraint. *Artificial Intelligence and Law*, 23(2):133–160, 2015.
111. E. L. Rissland. Example generation. In *Third Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, pages 280–288, 1980.
112. E. L. Rissland. Examples in legal reasoning: Legal hypotheticals. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 90–93, 1983.
113. E. L. Rissland. The ubiquitous dialectic. In *Proceedings of the Sixth European Conference on Artificial Intelligence*, pages 367–372, 1984.
114. E. L. Rissland. AI and legal reasoning. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence-Volume 2*, pages 1254–1260. Morgan Kaufmann Publishers Inc., 1985.
115. E. L. Rissland. Dimension-based analysis of hypotheticals from supreme court oral argument. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, pages 111–120. ACM, 1989.
116. E. L. Rissland and K. D. Ashley. A case-based system for trade secrets law. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 60–66. ACM, 1987.
117. E. L. Rissland and K. D. Ashley. A note on dimensions and factors. *Artificial Intelligence and Law*, 10(1-3):65–77, 2002.
118. E. L. Rissland and J. J. Daniels. A hybrid CBR-IR approach to legal information retrieval. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 52–61. ACM, 1995.
119. E. L. Rissland and J. J. Daniels. The synergistic application of CBR to IR. *Artificial Intelligence Review*, 10(5-6):441–475, 1996.
120. E. L. Rissland and D. B. Skalak. Combining case-based and rule-based reasoning: A heuristic approach. In *Proceedings of the 2nd International Joint Conference on Artificial Intelligence*, pages 524–530, 1989.
121. E. L. Rissland and D. B. Skalak. Interpreting statutory predicates. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, pages 46–53. ACM, 1989.
122. E. L. Rissland and D. B. Skalak. Cabaret: rule interpretation in a hybrid architecture. *International Journal of Man-Machine Studies*, 34(6):839–887, 1991.
123. E. L. Rissland, D. B. Skalak, and M. T. Friedman. Bankxx: a program to generate argument through Case-Base research. In *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, pages 117–124. ACM, 1993.
124. E. L. Rissland, D. B. Skalak, and M. T. Friedman. Bankxx: Supporting legal arguments through heuristic retrieval. *Artificial Intelligence and Law*, 4(1):1–71, 1996.
125. E. L. Rissland, D. B. Skalak, and M. T. Friedman. Evaluating a legal argument program: The bankxx experiments. *Artificial Intelligence and Law*, 5(1-2):1–74, 1997.
126. E. L. Rissland and E. Soloway. Overview of an example generation system. In *Proceedings of the 1st Annual National Conference on Artificial Intelligence*, pages 256–258, 1980.
127. E. L. Rissland, E. M. Valcarce, and K. D. Ashley. Explaining and arguing with examples. In *Proceedings of the Fourth AAAI Conference on Artificial Intelligence*, pages 288–294. AAAI Press, 1984.
128. E. L. Rissland and X. Xu. Catching gray cygnets: an initial exploration. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 151–160. ACM, 2011.
129. G. Sartor. Teleological arguments and theory-based dialectics. *Artificial Intelligence and Law*, 10(1-3):95–112, 2002.
130. G. Sartor. Doing justice to rights and values: teleological reasoning and proportionality. *Artificial Intelligence and Law*, 18(2):175–215, 2010.

131. R. Schank and R. Abelson. *Scripts, Plans, Goals and Understanding, An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum, Hillsdale, 1977.
132. 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.
133. D. B. Skalak and E. L. Rissland. Argument moves in a rule-guided domain. In *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, pages 1–11. ACM, 1991.
134. D. B. Skalak and E. L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial Intelligence and Law*, 1(1):3–44, 1992.
135. J. Smith and C. Deedman. The application of expert systems technology to case-based law. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, pages 84–93. ACM Press, New York, 1987.
136. S. T. Timmer, J.-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. ACM, 2015.
137. S. Toulmin. *The Uses of Argument*. Cambridge University Press: Cambridge, England, 1958.
138. B. Verheij. Arguments and defeat in argument-based nonmonotonic reasoning. In *Portuguese Conference on Artificial Intelligence*, pages 213–224. Springer, 1995.
139. B. Verheij. Formalizing value-guided argumentation for ethical systems design. *Artificial Intelligence and Law*, 24(4):387–407, 2016.
140. B. Verheij, J. C. Hage, and H. J. Van Den Herik. An integrated view on rules and principles. *Artificial Intelligence and Law*, 6(1):3–26, 1998.
141. D. Walton. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates: Mahwah, New Jersey, 1996.
142. D. Walton, C. Reed, and F. Macagno. *Argumentation schemes*. Cambridge University Press: Cambridge, England, 2008.
143. M. Wardeh, T. Bench-Capon, and F. Coenen. Padua: a protocol for argumentation dialogue using association rules. *Artificial Intelligence and Law*, 17(3):183–215, 2009.
144. A. Wyner and T. Bench-Capon. Argument schemes for legal case-based reasoning. In *Legal Knowledge and Information Systems - JURIX 2013: The Twentieth Annual Conference*, pages 139–149, 2007.
145. A. Wyner, T. Bench-Capon, and K. Atkinson. Towards formalising argumentation about legal cases. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 1–10. ACM, 2011.
146. L. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965.