

Case-Based Models of Legal Reasoning in a Civil Law Context

Kevin D. Ashley

Professor of Law and Intelligent Systems

University of Pittsburgh School of Law

Ashley@pitt.edu

1.0 Introduction

Researchers in Artificial Intelligence and Law (AI & Law) have developed a number of computational models of case-based legal reasoning. For example, see Hypo (Ashley, 1987; 1990), CABARET (Rissland & Skalak, 1991), GREBE (Branting, 1991; 1999), CATO (Aleven, 1997; 2003), BankXX (Rissland, Skalak, et al. 1996), and Split-Up (Zelevnikow, Stranieri, et al., 1995-1996). The models have originated in common law jurisdictions, often by lawyers/computer scientists influenced by common law legal practice and traditions. More recently, some of the case-based models have been developed by researchers in countries with more or less civil law traditions (Prakken & Sartor, 1997; Bench-Capon & Sartor, 2001).

The latter developments beg the question of whether AI & Law models of case-based legal reasoning are relevant in civil law jurisdictions. This paper addresses that and some related questions: To what extent is reasoning with precedents, so central in common law legal practice, practiced in civil law jurisdictions? Do civil law judges and practitioners reason with cases? If so, are there significant differences between the ways they reason with cases and those of their common law counterparts? What is the relevance in a civil law context of computational models of case-based and analogical legal reasoning developed in the common law tradition? Can/should they be incorporated into legal practice in civil law jurisdictions and how?

2.0 Case-based Analogical Reasoning in the Civil Law?

According to a recent comparative study, “precedent now plays a significant part in legal decision making and the development of law in all the countries and legal traditions that we have reviewed,” including nine civil law jurisdictions (e.g., Germany, France, Italy, and Spain) and two common law jurisdictions (the United Kingdom and New York State) (MacCormick and Summers, 1997, p. 531). This comprehensive, five-year study involved systematic research efforts to survey and relate uses of precedent in these jurisdictions. The study also concluded that, “all these systems accommodate justified legal change and evolution through judicial as well as legislative action, that is, through precedent.” (MacCormick and Summers, 1997, p. 535).

There are reasons to believe that at least some European legal systems are converging in their use of precedents: Given the “Europeanization of Europe” courts are beginning “to rely upon decisions not only of the European Court of Justice, but also of other Member State courts.” As judges in one European state confront “foreign values” from the others, they “need to question, and then to articulate, underlying assumptions.” As a result, “the style of opinion writing is becoming less ministerial, bold, and declaratory and more discursive, cautious, and fact-oriented. In short, judicial decisions are becoming more amenable to distinguishing and to ... use of the fact-based result of the decision in addition to the announced rationale and the discernable principles.” (Lundmark, 1998, pp. 223-4). Another important reason is that, “the proliferation of computers puts past decisions at the fingertips of judges and lawyers.” In addition, it is predicted that, as the “density of regulation” increases and as norms change more rapidly, “the same set of facts raises more legal issues than before. Consulting previous decisions (precedents) helps to chart one's way through the legal thicket of, for example, the burgeoning European private law.” Also, a “self-imposed adherence to precedent” will help judges “to reduce political disapproval, and to forestall legislative measures to restrict their ability to stray from precedent.” (Lundmark, 1998, pp. 223-4).

Differences across legal cultures in the uses of precedents, however, are subtle and profound. The above-mentioned study identified numerous substantial differences in legal reasoning with precedents across the civil law and common law jurisdictions studied (MacCormick & Summers, 1997, pp. 536-539):

1. **Few statements of facts:** “[M]ost officially published civil law opinions ... do not include ... detailed statements of facts.”...This matters because “what is reported substantially determines what is readily available to be used as a basis for argumentation in later cases.” (MacCormick and Summers, 1997, p. 536).
2. **Different significance:** In civil law judicial opinions “there is usually none of the detailed analysis and in-depth discussion of the point and purport of rulings on issues in prior cases.... [P]recedents are commonly conceived as loci of relatively abstract rules or (perhaps even more) principles, ... There is usually not, as in common law systems, a restriction of the binding element to a ruling on an issue of law considered in the special light of the material facts of the case. Thus, what we call the model of particular analogy plays far less part here.” (MacCormick and Summers, 1997, pp. 536-7).
3. **No focus on holdings:** In civil law systems, there is “no tradition of differentiating systematically in regard to a precedent opinion between *ratio decidendi* and *obiter dicta* – between holding and dictum – as in the common law.” (MacCormick and Summers, 1997, p. 537).
4. **Rules not contextualized:** “[R]ules in the common law are contextualized within and emerge from fact situations and fact patterns. ... [I]n most civil law systems ... the verbal formulations of general rules (statutory and other) and any relevant interpretive methodology are usually the primary determinants of their ultimate scope (always, of course, in conjunction with whatever article of statute or code may require interpretation in the decision).” (MacCormick and Summers, 1997, p. 537).
5. **No focus on distinguishing:** “[N]o sophisticated methodology of distinguishing precedents otherwise arguably applicable has developed in any of the civil law countries (again, constitutional cases aside), yet distinguishing has long been something of a high art among practitioners and judges in the common law countries.” In civil law countries, “tacit overruling or other departure” is employed. This obscures lines “that ought to be drawn between closely analogical precedents that point in different directions.” (MacCormick and Summers, 1997, pp. 538-9).
6. **Lines of precedents required:** “[I]n most of the civil law countries, a single precedent is usually not on its own sufficient to count as authoritatively settling a point of law (again, constitutional cases aside). Several precedents, that is a ‘line’ of precedents, are usually required....” (MacCormick and Summers, 1997, p. 538).
7. **Subsequent court departures:** “[A] vital difference concerns the liberty of even lower courts to depart from a single higher-court precedent, or even

from a line of several precedents....In Italy, Germany, Finland, France and Spain at least, apparently settled points can be reopened even by trial courts of general jurisdiction on their own judgment as to what is the law, or good law.” (MacCormick and Summers, 1997, p. 538).

8. **Tacit following, tacit departures:** “Precedents may be followed, confirmed even, by courts of final instance without express citation or mention.” Likewise, “[I]n five of the civil law systems in our study, Sweden, Italy, Spain, France, and Norway, the higher and highest courts consciously, and with some regularity, depart from precedent without even mentioning this fact.” (MacCormick and Summers, 1997, p. 539).
9. **Not formal sources of law:** “[T]hese features ... are symptomatic of a conception of precedent that deems it something other than or less than a full-dress formal source of law and which, accordingly, has somewhat lower normative force.” (MacCormick and Summers, 1997, p. 539).

Even when courts in civil law and common law systems are attempting to achieve uniform application of the very same law, such as the U.N. Convention on the International Sale of Goods (CISG), fundamental differences may be evident in “what the national courts consider to be primary and secondary sources of legal authority” and “differences across legal cultures in the understandings even of what a judicial decision is”. (Curran, 2001, pp. 67f).

[W]here a U.S. judge striving to apply the CISG uniformly would be prepared to consult prior CISG case law, a French judge would expect to consult scholarly commentary rather than the judicial decisions themselves. Moreover, a U.S. judge would be perplexed by a French judicial application of the CISG, because the French court opinion might well consist of one sentence without any clear description of the case's underlying factual scenario, and essentially be inaccessible without the explanatory scholarly commentary that French lawyers seek when trying to understand French judicial decisions.

Conversely, a French judge assessing United States CISG case law instinctively would look for *la doctrine*, the scholarly commentary that occupies a privileged position of influence on French court adjudications, but which, to a common-law trained legal mind, may be perceived as tainted by the scholar's interpretive subjectivity, not to speak of by the lowly status of American scholars in terms of their influence on court decisions. (Curran, 2001, p. 68).

Civil law jurists and legal practitioners, of course, must decide for themselves the utility of reasoning with precedents and the likelihood that it will occur in the foreseeable future. The above observations, however, suggest two

alternative possible states of affairs that may evolve if a traditionally civil law jurisdiction should come to use precedents in legal reasoning: the Abstract Precedent Scenario or the Fact-Based Precedent Scenario:

Abstract Precedent Scenario: In this scenario, a precedent is deemed useful, if at all, as an indication that another or higher court has referred to an abstract rule or principle in connection with the particular article of a statute or code that requires interpretation in the current problem and/or has formulated the abstract rule or principle in a particular way. The precedent contains little if any description of the facts in which the abstract rule or principle was applied. This is of little concern because those facts are of no particular interest to the subsequent court and bear little relevance to the use it will make of the precedent.

Fact-Based Precedent Scenario: In this scenario, by contrast, a precedent is useful as an indication that another or higher court has come to a particular decision in the context of a fact situation relevantly similar to that of the problem. The precedent contains a rich description of the facts involved. The decision may involve any or all of the following – that under factually similar circumstances, the plaintiff won/lost a particular:

- a) kind of legal claim,
- b) issue involved in that kind of legal claim, or
- c) issue involved in that kind of legal claim for a particular reason.

Clearly, common law uses of precedents involve drawing legal inferences from a comparison of current problems and past cases on their facts. In this sense, the Fact-Based Precedent is much closer to common law uses of precedents. In formulating the description of the Fact-Based Precedent scenario, however, I have deliberately avoided making assumptions about why the precedent is useful. In particular, I assume that it may be useful even if the jurisdiction is not like a common law jurisdiction in that it does *not* adhere to the doctrine of *stare decisis* that similar cases should be decided alike. I assume that the precedent may simply be a more or less influential example. By observing

how the prior court decided a legal dispute involving similar factual circumstances, a subsequent court may simply be reminded of what claims, issues or reasons are relevant in that type of factual scenario. In addition, it may be persuaded that a similar decision in the problem is a good result in a normative sense. After comparing the facts of the cases, it may even be persuaded that the reasons for the decision in the prior case do not apply in the current problem and that a different result would be better, normatively. Notice that any of the above may be true, and the precedent may thus be potentially useful, even if it is not the case that the subsequent court is bound by *stare decisis* or by the rules of a hierarchical court system to follow the decisions of a prior or higher court.

The essential difference between the Abstract and Fact-Based Precedent Scenarios is that the latter emphasizes the importance of comparing the facts of the current problem with the factual scenario in the precedent. This assumes, of course, that the opinion in the precedent reports the factual scenario. It may even be the case that the Fact-Based Precedent is useful primarily for the same reason as the Abstract Precedent, that is, because it refers to or provides a formulation of an “abstract rule or principle in connection with the particular article of a statute or code that requires interpretation in the current problem.” The Fact-Based Precedent, however, will also be useful because it provides an example of a court’s application of the abstract rule or principle in a factual context which can be compared with the facts of the problem. In other words, the factual context of the precedent and the prior court’s decision are important components of the Fact-Based Precedent’s significance; the factual context and decision help to demonstrate what the rule or principle (and thus the statute) means.

Presumably, the states of affairs concerning the use of precedent in a civil jurisdiction now and in the future lie somewhere between these two descriptions. For purposes of this paper, the important point is that the computational models of case-based legal reasoning that have been developed in AI & Law have all been designed to model legal inferences from the kind of fact-oriented case comparisons that underlie the Fact-Based Precedent approach. These

computational models are not necessarily models of precedent; they do not necessarily assume that the doctrine of *stare decisis* is followed. They do, however, model arguments that a problem should be decided in the same way as or differently from a case or cases based on drawing factual analogies or distinctions. In the language of the McCormick and Summers study, they provide computational implementations of a “model of particular analogy.” To the extent that the Fact-Based Precedent approach is irrelevant in a jurisdiction, then so will be these AI & Law models of case-based legal reasoning.

3.0 Uses of Precedent and Automated Legal Information Retrieval

Current automated legal information retrieval systems can help practitioners retrieve precedents under either state of affairs: the Abstract Precedent Scenario or the Fact-Based Precedent Scenario. These database systems, like Westlaw in the United States, comprise all of the published judicial opinions issued by courts in a particular jurisdiction. The opinion texts are processed to transform the words to remove endings (i.e., stemming), to remove stop words (i.e., words like “a”, “an”, and “the” so common that they may be of little value in retrieval) and to identify various features, such as, citations to statutory or constitutional provisions or to previous cases, significant phrases and special indexing concepts. The document is then indexed in an inverted index by each remaining word and the other features. Using the index, a system can retrieve all documents that contain a particular word (i.e., terms, phrases, citations, or concepts) or a particular set of features.




Figure 1: Bayesian Inference Network Retrieval Model

In a modern full-text legal information system, the database is constructed as a Bayesian inference network, a mechanism for representing conditional probabilities and drawing inferences from them. In Figure 1, the top half represents the documents d_i in the system's database, indexed by their features f_i . The bottom half represents a query Q , presented to the system and processed much like a short document into terms, phrases, citations, and concepts. The fact that a query Q has been observed with certain features f_i is treated as some evidence that a particular document d_1 satisfies that query. It is also some evidence that d_2 satisfies, d_3 , etc. The Bayesian inference network's task is to determine for each document how much evidence the query provides and to rank the documents accordingly. It estimates the probability that a particular document d_i satisfies query Q using TF/IDF values. These values depend on the frequency of the term or other feature in the document (TF) and in the collection as a whole (IDF). They increase with TF and decrease with the IDF. Thus, a citation that appears frequently in the document but rarely in the corpus leads to a high estimated probability that the document satisfies a query that also includes that citation. The system ranks this and other documents according to the magnitude

of the probabilities and presents the n top-ranking documents to the user (Turtle, 1995, p. 33).

It will be apparent that such a full-text legal information system can assist practitioners in retrieving cases in either the Abstract Precedent Scenario or Fact-Based Precedent Scenario.

To the extent that practitioners in a legal system are concerned only with retrieving Abstract Precedents, then legal opinions will likely continue to contain statements of abstract rules or principles and citations to particular articles of a statute or code to which they pertain, but will not contain detailed descriptions of the case's factual context. Such fact-deficient opinions, nevertheless, can be processed and retrieved with systems that employ an inverted index and a Bayesian inference network.

Retrieval of Abstract Precedents might work even better, however, if the opinions also contained descriptions of the facts of a case. Even if the goal in retrieving an Abstract Precedent is only to identify the "abstract rule or principle in connection with the particular article of a statute or code that requires interpretation in the current problem" and even if a statement of the facts of the prior case is deemed to be irrelevant to the interpretation of the abstract rule or principle, there still would be some utility if judicial opinions included more extensive descriptions of factual contexts. For one thing, a practitioner may not be sure of the terms employed in the abstract rules or principles. Full-text legal information retrieval tools make it considerably easier to retrieve past cases that satisfy queries described either in terms of the legal concepts involved in the abstract rules or principles, or of citations to relevant statutory provisions, or simply of the problem's relevant facts. Moreover, the terms of the abstract rules or principles are ambiguous. The same terms may appear in a variety of contexts. By inclusion of some of the factual circumstances in the problem, retrieving factually similar cases would increase the chances that the abstract rules or principles retrieved are the ones relevant in the user's intended context.

To the extent that practitioners are interested in using full-text legal information systems to retrieve Fact-Based Precedents, of course, it is then

essential that the stored opinions and the query contain a description of the facts. Systems like Westlaw can accept fairly detailed natural language descriptions of a problem's facts, turn them into queries, retrieve and rank documents using the Bayesian inference network, and present the user with precedents involving remarkably similar factual scenarios. It outputs lists of documents ranked in terms of probabilities that the documents satisfy the query, highlights the query's terms in the text, and even highlights parts of the texts where those terms are highly concentrated. It also makes it trivially easy to retrieve any statutes or cases cited in a retrieved text or any cases that cite the retrieved text.

Indeed, systems like Westlaw have remarkable strengths. It is easy to input queries as text. If the judicial opinions are already in electronic files, the process of adding them to the database is completely automatic! New judicial opinions can be speedily processed, added to the database and indexed in the inverted index without human intervention. As a result, the databases can be enormous and comprehensive. Determining how relevant documents are to a query (i.e., assigning the TF/IDF values and calculating the probabilities) is performed automatically. In addition, the Bayesian inference network provides a simple and effective way of combining evidence from multiple document representation schemes (e.g., terms, citations, phrases, and other indexing concepts.) This is important, as Westlaw can also factor in its enormous manually-prepared subject-matter indexes.

In sum, whether practitioners are interested in retrieving Abstract Precedents or Fact-Based Precedents, a Westlaw-type full-text legal information system would serve well, especially if judges are prepared to include in their opinions more elaborate statements of case facts.

There are, however, things that a full-text legal information system cannot do. It cannot draw legal inferences from a comparison of the facts of problem and cases. Nor can it flag what facts are important in that comparison. Its TF/IDF-based relevance measure does not relate especially well to legal concepts of relevance or to the ways that practitioners would use the cases in argumentation (at least in common law styles of argumentation). Its outputs of ranked

documents with highlighted query terms, while very useful, are not as helpful as they might be. The systems cannot interpret the texts, even to identify who won, on what claims, involving what facts. Finally, they cannot predict the outcomes of a problem.

All of these tasks are beyond the capabilities of current full-text legal information systems. Achieving them has been the goal of extensive research in AI & Law, and in particular, of the computational models of case-based legal reasoning described below.

4.0 Computational Models of Case-Based Legal Reasoning

A computational model of case-based legal reasoning requires (1) a scheme for representing the facts of cases and problems that are legally significant and why, (2) a means for assessing the relevance of cases to a problem, and (3) a mechanism for comparing cases and drawing legal inferences.

At least two representational schemes have been developed in AI & Law.¹ The first is based on Dimensions (Ashley, 1987; 1990) or their simpler relatives Factors (Aleven, 1997; Aleven & Ashley, 1994); the second involves Exemplar-Based Explanations (Branting, 1991; 1999). Dimensions capture stereotypical patterns of fact that tend to strengthen or weaken a side's position on a claim. EBEs capture an explanation of how the legal conclusions are justified in terms of the facts.

Both schemes support relevance measures that directly relate to the ways in which cases are used in legal arguments. Both support retrieving relevant cases from a database, comparing problems to cases, drawing legal inferences, and explaining them in well-formed legal arguments. In this respect, these computational models have a big advantage over full-text legal information retrieval systems. Westlaw cannot draw legal inferences from the cases it retrieves nor show how they can be used in arguments.

¹ Approaches integrating rules and either Dimensions or Factors include (Rissland & Skalak, 1996; Prakken & Sartor, 1997; Bench-Capon & Sartor, 2001). A hybrid connectionist approach to some kinds of case-based legal decisions has been implemented in (Zeleznikow, Stranieri, et al., 1995-1996).

On the other hand, in constructing the databases of cases using either Dimensions or EBEs, it was necessary for someone to read the case opinions and manually fill-in the case representations. Problems submitted for analysis must be similarly represented. In this respect, the AI & Law models are at a disadvantage compared to, say, Westlaw. They cannot read and understand the opinion texts any better than Westlaw can. But unlike Westlaw, in order for the AI & Law models to process the cases for purposes of retrieval, inference and explanation, substantive aspects of the cases have to be represented. Since the special representation schemes make it harder to input new documents and queries, the AI & Law models employ much smaller databases of, at most, a few hundred cases.

4.1 Dimension- and Factor-based Representation Schemes

Dimensions have been implemented for the domain of trade secret misappropriation law, among others. In the US, trade secret law protects developers of secret information that confers a competitive advantage from competitors who gain and use the information through a breach of a confidential relationship or by improper means. In the US, trade secret law is mainly state law, either common law or statutory. The main sources of trade secret law are the Restatement First of Torts, Section 757, a scholarly restatement cited by and relied upon in state court opinions, the Uniform Trade Secret Act, and the Restatement Third of Unfair Competition, Sections 39 - 45.

Even where the claim is statutory, the statutes are not comprehensive codes. For instance, the UTSA definition of a trade secret is: “Trade secret’ means information, ... that: (i) derives independent economic value, ... from not being generally known to, and not being readily ascertainable by proper means ... and (ii) is the subject of efforts that are reasonable under the circumstances to maintain its secrecy.” The Restatement First states, “One ... is liable [for trade secret misappropriation if] (a) he discovered the secret by improper means, or (b) his disclosure or use constitutes a breach of confidence” It relegates the “definition” to Comment b: “A trade secret may consist of any formula, pattern, device or compilation of information which is used in one's business, and which

gives him an opportunity to obtain an advantage over competitors who do not know or use it.” Comment b goes on to say, “An exact definition of a trade secret is not possible. 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.”

Dimensions in the HYPO program expand upon that list of factors. Some examples include the Dimensions in Figure 2.

Security-Measures: plaintiff's claim is stronger the more security measures it took to protect info.

Disclosure-In-Negotiations: plaintiff's claim is stronger to the extent it did not disclose the secret to defendant in negotiations.

Agreed-Not-To-Disclose: plaintiff's claim is stronger to the extent it entered into a nondisclosure agreement with the defendant.

Employee-Sole-Developer: plaintiff's claim is stronger to the extent that defendant was not the sole developer of the information.

Secrets-Disclosed-Outsiders: plaintiff's claim is stronger the fewer disclosures of information were made to outsiders.

Outsider-Disclosures-Restricted: plaintiff's claim is stronger to the extent that disclosees were restricted from disclosing the information to others.

Competitive-Advantage: plaintiff's claim is stronger the greater competitive advantage defendant gained by access to plaintiff's information.

Bribe-Employee: plaintiff's claim is stronger the more money, stock, or other benefits the defendant gave to plaintiff's former employees to switch employment.

Brought-Tools: plaintiff's claim is stronger to the extent the former employee brought product-related tools to defendant.

Figure 2: Sample Dimensions in HYPO

HYPO's Dimensions are a kind of expert knowledge. Each one relates to a stereotypical fact pattern identified by legal scholars in the field. For each one, there is at least one case where a judge said that the underlying pattern strengthened or weakened a claim.

HYPO's Dimensions were highly structured objects, complete with preconditions that determined when they applied and ranges of possible values that indicated how extreme an example of the Dimension a case presented. For instance, Security-Measures' range comprised sets of eight types of security measures commonly taken. A case's value on this Dimension could range from the empty set, the weakest value for plaintiff, to the set of all possible measures, the strongest value for plaintiff. Other Dimensions had binary ranges, such as whether or not a plaintiff had disclosed secrets to defendant in negotiations. In order to support experimentation with an intelligent tutoring system to teach students to make case-based arguments, work on the CATO program introduced Factors as a means of simplifying Dimensions. Factors are binary; a Factor's value is true if it applies to the facts of a case and false if it either does not apply or it is not known whether it applies. Thus, if the Factor, Security-Measures, applies in a case it represents a strength for plaintiff regardless of what non-empty set of measures were taken. Otherwise, the Factor does not apply in a case, that is, either it is known not to apply or it is unknown whether it applies.

For instance, consider the following sample fact situation based on a real trade secret law case, *Mason v. Jack Daniels Distillery*.

In 1980, a restaurant owner named Mason developed a combination of Jack Daniel's whiskey, Triple Sec, sweet and sour mix, and 7-Up to ease a sore throat. He promoted the drink, dubbed "Lynchburg Lemonade" for his restaurant, "Tony Mason's, Huntsville", served it in Mason jars and sold T-shirts. Mason told the recipe only to his bartenders and instructed them not to reveal the recipe to others. The drink was only mixed out of the customers' view. [F6 Security-Measures (p)] The drink comprised about one third of the sales of alcoholic drinks. Despite its extreme popularity, no other establishments had duplicated the drink, but experts claimed it could easily be duplicated. [F15 Unique-Product (p); F16 Info-Reverse-Engineerable (d)] In 1982, Randle, a sales representative of the Jack Daniel's Distillery, visited Mason's restaurant and drank Lynchburg Lemonade. Mason disclosed part of the recipe to Randle in exchange, Mason claimed, for a promise that Mason and his band would be used in a sales promotion. [F1 Disclosure-in-Negotiations (d)] Randle recalled having been under the impression that Mason's recipe was a "secret formula". [F21 Knew-Info-Confidential (p)] Randle informed his superiors of the recipe and the drink's popularity. A year later, the Distillery began using the recipe to promote the drink in a national sales campaign. Mason did not participate in the promotion or receive other compensation.

Figure 3: Facts of the *Mason* Case

Experts in trade secret law would recognize five stereotypical fact patterns that strengthen or weaken the plaintiff Mason's trade secret claim against defendant Jack Daniel Distillery. Each corresponds to a Factor and has been inserted into the above text, along with an indication of which side it favors, immediately after the sentence that justifies its application. Thus, Factors F6, Security-Measures, F15, Unique-Product, and F21 Knew-Info-Confidential all favor the plaintiff (p). Factors F16, Info-Reverse-Engineerable and F1, Disclosure-in-Negotiations also apply but favor the defendant (d).

4.2 Issue-Based Prediction of Problem Outcomes

Of course, little can be determined about the likely outcome of the *Mason* case, or about what reasonable arguments can be made for or against plaintiff's claim solely from the information that these Factors are present, that they compete, or even that three Factors favor plaintiff and two favor defendant.

With a Domain Model that relates the Factors to issues in trade secret law and a database of trade secret cases represented in terms of, and indexed by, Factors, however, a program can frame and test hypotheses about which side is likely to win, explain its predictions, and even make the strongest arguments for and against each side. We developed the Issue-Based Prediction (IBP) program for this purpose (Brüninghaus & Ashley, 2003).




Figure 4: IBP's Domain Model

IBP's Domain Model, shown in Figure 4, is based on the Restatement First of Torts, Section 757, and on the Uniform Trade Secrets Act. It identifies two main issues and five sub-issues involved in a claim of trade secret misappropriation, and it provides a logical framework for these issues. Plaintiff must show that the information is a trade secret and was misappropriated. It can show the former by showing that the information is valuable and that it took efforts to maintain secrecy. It can show that the information was misappropriated by showing either that the information was obtained through improper means or that it was used in breach of a confidential relationship. None of the sub-issues is defined in logical terms. Instead, each is related to a set of Factors. For each such Factor, the legal reason why the Factor is important is that it is relevant to the sub-issue(s). CABARET was the first program to represent legal predicates in terms of factors and cases for purposes of argument-making (Rissland & Skalak, 1991); IBP does so for the purpose of predicting outcomes (Brúninghaus & Ashley, 2003).

Given a new problem situation, represented as a set of Factors, IBP uses the Domain Model to identify the issues relevant in the problem. For each issue, it determines if the issue-related Factors all favor the same side. If so, it predicts that side will win the issue. If, however, the issue-related Factors favor conflicting parties, IBP retrieves cases from the database that share those Factors and examines their outcomes. It poses a hypothesis that the side should win corresponding to the winner of the majority of the retrieved cases. It then tests the hypothesis against the retrieved cases. If there are no counterexamples (i.e., no cases won by the other side), the hypothesis is confirmed; IBP predicts that side will win the issue. If there are counterexamples, IBP determines whether the counterexamples can be explained away.

In explaining away counterexamples, IBP attempts to distinguish them from the problem situation. As indicated in the MacCormick and Summers study, distinguishing a cited case is an important task in common law legal argument. It means finding legal reasons that explain the result in the cited case but that do not apply in the problem. In IBP (and also in CATO and HYPO), the legal reasons are associated with Factors that favor the result in the cited case, but which are not present in the problem, or Factors that favor the opposite result in the problem, not present in the cited case. For example, the counterexample may have had some particularly strong Factor favoring the opposing side that explains why that side won and that is not present in the problem. Such Factors are called “knock out” Factors or KO-Factors, for short.

If all of the counterexamples can be explained away, the program predicts the majority side should win the issue. If not, it abstains from a prediction on that issue. If the hypothesis is too specific to retrieve any cases, IBP broadens the query by relaxing the constraints systematically in search of a hypothesis for which case examples can be found and from which the more specific but untestable hypotheses would follow *a fortiori*. After addressing each relevant issue, IBP employs its Domain Model to make an overall prediction or abstain.

For the *Mason* problem, as shown in Figure 5, IBP identifies three relevant issues, Security-Measures, Confidential-Relationship, and Info-Valuable, predicts

that plaintiff will win each one and the overall claim for trade secret misappropriation. For the latter two issues, IBP finds conflicting issue-related Factors and conflicting cases, so it engages in hypothesis-testing (or theory-testing.) Since for each issue, plaintiff won the majority of cases, IBP hypothesizes that plaintiff will win the issues in *Mason*. It then attempts to explain away any counterexamples. Here, it successfully explains away the *Ecologix* and *National Rejectors* counterexamples as involving KO-Factors not present in the *Mason* facts.

```

Prediction for MASON, which was won by ???
  Factors favoring plaintiff: (F21 F15 F6)
  Factors favoring defendant: (F16 F1)

Issue raised in this case is SECURITY-
MEASURES
  Relevant factors in case: F6(P)
PLAINTIFF.

Issue raised in this case is CONFIDENTIAL-
RELATIONSHIP
Relevant factors in case: F1(D) F21(P)
Theory testing has no clear outcome, try to
explain away exceptions.
Cases won by plaintiff:
  BOEING (F1 F4 F6 F10 F12 F14 F21)
  BRYCE (F1 F4 F6 F18 F21)
  DEN-TAL-EZ (F1 F4 F6 F21 F26)
  DIGITAL-DEVELOPMENT (F1 F6 F8 F15 F18
F21)
  FOREST-LABORATORIES (F1 F6 F15 F21)
  GOLDBERG (F1 F10 F21 F27)
  LASER (F1 F6 F10 F12 F21)
  LEWIS (F1 F8 F21)
  NATIONAL-INSTRUMENT (F1 F18 F21)
  VALCO-CINCLINNATI (F1 F6 F15 F21 F10 F12)
Cases won by defendant:
  ECOLOGIX (F1 F19 F21 F23)
Trying to explain away the exceptions
favoring DEFENDANT
  ECOLOGIX can be explained away because of
the unshared ko-factor(s) (F23 F19).
Therefore, PLAINTIFF is favored.

Issue raised in this case is INFO-VALUABLE
  Relevant factors in case: F16(D) F15(P)
Theory testing has no clear outcome, try to
explain away exceptions.
Cases won by plaintiff:
  AMERICAN-CAN (F4 F6 F15 F16 F18)
  HENRY-HOPE (F4 F6 F15 F16)
  ILG-INDUSTRIES (F7 F10 F12 F15 F16 F21)
  KAMIN (F1 F10 F16 F18 F15)
  KG (F6 F14 F15 F16 F18 F21 F25)
  KUBIK (F7 F15 F16 F18 F21)
  TELEVATION (F6 F10 F12 F15 F16 F18 F21)
Cases won by defendant:
  NATIONAL-REJECTORS (F7 F10 F15 F16 F18
F19 F27) Trying to explain away the
exceptions favoring DEFENDANT:
  NATIONAL-REJECTORS can be explained away
because of unshared ko-factor(s) (F27
F19). Therefore, PLAINTIFF is favored.

Outcome of the issue-based analysis:
  For issue CONFIDENTIAL-RELATIONSHIP,
PLAINTIFF
  is favored.
  For issue SECURITY-MEASURES, PLAINTIFF
is
  favored.
  For issue INFO-VALUABLE, PLAINTIFF is
favored.
=> Predicted outcome for MASON is
PLAINTIFF, which
  is correct.

```

Figure 5: IBP's Output for *Mason*

4.3 Case-Based Arguments with Factors, Issues, Cases

Does this mean that defendant in *Mason* is doomed? Not necessarily. IBP's prediction for a problem is an empirical prediction based on the cases in its database. It may be wrong. The defendant may win, and, in any event, may be able to make good legal arguments. The CATO program can find in the database the least distinguishable, most relevant cases the defendant can cite without fear of plaintiff's responding with a more relevant pro-plaintiff counterexample. It can

use them to make arguments why defendant in *Mason* should win despite the predictions.

In CATO, the basic measure of relevance is on-pointness; a case is on point if it shares at least one Factor with the problem. One case is more on point than another case if the second case's set of Factors shared with the problem is a subset of those shared by the first case and the problem. As HYPO before it, CATO partially orders all of the relevant cases in terms of their on-pointness to the problem in a data structure called a Claim Lattice (Ashley, 1987; 1990). Cases along a branch of the Claim Lattice that are closer to the root node, representing the problem's set of applicable factors, are more on point than those farther down a branch. For instance, in analyzing *Mason*, suppose CATO retrieved from the database a (hypothetical) on-point case *c*, represented as having been won by defendant and with the following factors: F27(d) Disclosure-In-Public-Forum, F19(d) No-Security-Measures, F18(p) Identical-Products, F16(d) Info-Reverse-Engineerable, F4(p) Agreed-Not-To-Disclose, F1(d) Disclosure-In-Negotiations. Moreover, suppose CATO finds no pro-plaintiff case more on-point (i.e., closer to the root of the Claim Lattice) than *c*; in other words, if defendant argues that defendant should win in *Mason*, citing *c*, plaintiff could not cite any more on-point pro-plaintiff case as a counterexample with which to trump defendant's argument. CATO's arguments comparing the *Mason* problem and case *c* are shown in Figure 6. (For simplicity, this is actually a composite of two arguments.)

==> Point for Defendant as Side-1:

WHERE: Plaintiff's product information could be learned by reverse-engineering (F16) and plaintiff disclosed its product information in negotiations with defendant (F1)

DEFENDANT should win a claim for Trade Secrets Misappropriation.

CITE: Case *c*

<== Response for Plaintiff as Side-2:

Case *c* is distinguishable, because: In *c*, plaintiff disclosed its information in a public forum (F27). Not so in *Mason*. In *c*, plaintiff did not adopt any security measures (F19).

Not so in *Mason*. In *Mason*, plaintiff adopted security measures (F6). Not so in *c*. In *Mason*, plaintiff was the only manufacturer making the product F(15). Not so in *c*. In *Mason*, defendant knew that plaintiff's information was confidential F(21). Not so in *c*.

COUNTEREXAMPLES: None.

==> Rebuttal for Defendant as Side-1:

In *c*, plaintiff did not adopt any security measures (F19). This was not so in *Mason*. However, this does not amount to an important distinction. In *Mason*, plaintiff disclosed its product information in negotiations with defendant. In both cases, therefore, plaintiff showed a lack of interest in maintaining the secrecy of its information.

In *c*, plaintiff disclosed its information in a public forum (F27). This was not so in *Mason*. This however is not a major distinction. First, in *Mason*, plaintiff disclosed its information to defendant during negotiations and plaintiff's information could be discovered by reverse engineering plaintiff's product. It follows that in both cases, defendant obtained or could have obtained its information by legitimate means. Second, in *Mason*, plaintiff conveyed its information to defendant in the course of negotiations. In both cases, therefore, plaintiff showed a lack of interest in maintaining the secrecy of its information.

Figure 6: CATO's Best Argument for Defendant in *Mason*

As shown in Figure 6, particularly the defendant's rebuttal, CATO argues from a more general normative viewpoint that the two cases are fundamentally similar and should be decided alike. Using one set of argument evaluation criteria, CATO does not deem plaintiff's response distinguishing the case *c* as particularly successful, even though that case has two strong pro-d Factors, F19 and F27 not shared in *Mason*. In the rebuttal it finds that defendant can downplay these distinctions, arguing that they do not make *c* significantly worse for the plaintiff than the situation in *Mason*, and, therefore, that *Mason*, like *c*, should be decided for the defendant.

In making determinations about the significance of distinctions and whether they can be emphasized or downplayed, CATO employs a different knowledge representation structure, the Factor Hierarchy. For every factor, the Factor Hierarchy relates it to legal reasons why it matters in terms of the higher level issues of trade secret law. CATO draws on this information in constructing the rebuttal arguments of Figure 6 (Aleven, 1997).

Both IBP and CATO have been evaluated empirically. In an experiment, IBP outperformed a variety of other algorithms in predicting the outcomes of cases, achieving an accuracy of 91.4%. A naïve Bayes approach came in second with 86.5% accuracy, but it, unlike IBP cannot generate explanations of its predictions. IBP's Domain Model and its database of cases represented in terms of Factors enable it to formulate and test hypotheses about which side should win, and to evaluate the hypotheses using techniques for distinguishing and explaining away counterexamples. CATO has been implemented as a tutoring system for teaching law students to make legal arguments, and its pedagogical benefits have been demonstrated empirically in research evaluations. CATO's argumentation measures can also be used to base predictions on the best cases (e.g., least distinguishable untrumped counterexamples). This method yielded an accuracy of 77.8% (Brüninghaus & Ashley, 2003; Alevén, 2003).

Beyond the domain of trade secret misappropriation, factor-based AI & Law models have been applied to particular issues in personal income tax law and bankruptcy (Rissland & Skalak, 1991; Rissland, Skalak, et al. 1996)

4.4 Alternative Representation Scheme: EBEs

As noted above, another scheme has also proven useful for representing relevant facts in a case, EBEs, developed in connection with the GREBE program (Branting, 1991; 1999). EBEs represent not only the relevant facts of a case but also aspects of the judge's analysis of their legal significance in justifying her decision. In GREBE, EBEs were applied to represent workman's compensation cases, a statutory domain.

While a detailed discussion of EBEs and their use is beyond the scope of this paper, it is interesting to contrast the approach with that in CATO and IBP. The EBE representation requires identifying the statutory terms that are disputed in a case, and for each one, representing a brief explanation of why the judge decided that the term was (or was not) satisfied in the case. The explanation includes the "criterial" facts, the particular facts that the judge deemed legally significant in his decision regarding the applicability of that term in the case.

These facts are expressed in a relational language and linked to the appropriate statutory terms in a semantic network. In turn, those terms are related to other statutory terms through a logical structure that represents the judge's logical path through the statutory rules to a conclusion. The cases are stored in a database, indexed by the statutory terms of which they are positive or negative instances. Given a new case, GREBE recursively attempts to apply the statutory rules to the facts, and where particular statutory terms are not further defined by rules, it retrieves cases indexed by those terms and attempts to map the criterial facts, and accompanying explanation, from the case onto the problem's facts. GREBE measures relevance between a problem and a retrieved case as the fraction of the number of unshared and shared criterial facts between them. It selects the best-matching cases and generates a legal argument by analogy, elaborating the criterial facts shared by the problem and cited cases. Likewise, it can distinguish a cited case from a problem in terms of unshared criterial facts. GREBE did not generate predictions of who would win in a problem; it presented its arguments by analogy characterizing them as stronger or weaker depending on the matching of criterial facts.

The EBE scheme puts a premium on consistently representing the corresponding parts of cases' and problems' semantic networks, so that the approach of matching criterial facts will work. This is not easy given the vast number of ways to express any such explanation and the difficulty of determining exactly what a judge's rationale is and at what level of abstraction to express it.

Given the problems of manually constructing consistent EBEs across many cases, it is interesting that IBP was able to make predictions with 91.4% accuracy even though it does not have a representation of the judge's actual analysis or rationale for any case, only the cases' basic facts. Using its Domain Model, IBP can generate reasonable interpretations of how a court might analyze a particular issue given a problem's facts, and that proved enough to enable it to do a good job of formulating and testing prediction hypotheses based on past cases.

4.5 Connecting to Case Texts and Full-text Legal IR

The practical promise of factor-based and other AI & Law approaches depends on the extent to which they can help deal with intelligently processing cases – and case texts – on a much larger scale. Extending CATO to other legal domains would be greatly facilitated if techniques were available for semi-automatically indexing cases by their applicable Factors. As a practical matter, applying IBP or CATO to assist legal practitioners in predicting outcomes of real problems and generating alternative arguments depends on the extent to which they can be integrated with and add value to full-text legal information services like Westlaw.

There are at least three ways to pursue the goals of connecting the AI & Law models with case texts and full-text legal information retrieval.

First, an AI & Law program can be used to seed inquiries to a legal IR system like Westlaw. An integrated IBP/CATO program could be implemented for other legal domains beside trade secret law. This would largely be a matter of developing a Domain Model, Factors, Factor Hierarchy and cases for each new domain. (It may be possible to integrate the Domain Model and Factor Hierarchy into one model.) Each such specialized domain coverage might involve tens or hundreds of cases. Legal researchers would use the IBP/CATO program to research problems in the specialized area, generating predictions and arguments as needed. To the extent they liked the cases they found with IBP/CATO, they would use them to “seed” and launch queries into Westlaw for additional cases. For instance, if the user were interested in case *c* above, it is a trivial matter to retrieve all cases it cites or that cite it using the KeyCite or Shepard’s citation services available through Westlaw or Lexis. If the user were interested in cases like *c*, with factors F27, F19, F18, F16, F4, F1, English titles or descriptive phrases associated with those Factors could be fashioned automatically into a natural language query to Westlaw. Informal experience with such queries indicates a reasonably good chance that the cases retrieved by Westlaw will include some that are trade secret cases involving the relevant fact patterns. Of course, a reader must manually read the returned cases to be sure.

Second, an automated approach to the seeding of such queries that also highlights relevant portions of the retrieved case texts has been developed in the SPIRE program (Risland & Daniels, 1996). The program has a database of cases dealing with the issue of whether a bankruptcy plan has been submitted in good faith. The cases are represented in terms of features not unlike Dimensions. Given a new problem represented as a set of such features, SPIRE retrieves relevant cases, organizes them into a Claim Lattice, and selects the most on point cases. It then passes the texts of the selected cases to the relevance feedback module of INQUERY (Callan, et al., 1992), a full-text information retrieval system with a database of legal texts. These case texts seed a query, in effect, instructing INQUERY to retrieve more texts like these.

In experiments (Risland & Daniels, 1996). SPIRE found new and important cases very similar to the inputted problems (i.e., involving the same kind of legal stories), thus offering the possibility of semi-automating the maintenance of an AI & Law model's case database directly from full-text legal information retrieval systems. A SPIRE user can even indicate the particular features of interest, and the program will automatically highlight the parts of the texts of the retrieved cases that correspond to that feature. The highlighting mechanism works on the same principle. The program has a database of short passages for each feature. It assembles the passages associated with the feature of interest into a query submitted to INQUERY's relevance feedback module. That program, now using the texts of all the retrieved cases as its database, pulls up and highlights the passages in the cases most similar to the query.

Third, techniques for automatically extracting Factor-related information from textual cases for purposes of automatic highlighting and indexing, are also under development. SMILE (for SMart Index LEarner) employs a combination of information extraction tools and machine learning. Using the ID3 learning algorithm, SMILE learns decision trees for classifying sentences as positive or negative instances of a Factor. The positive instances are sentences in textual summaries of case opinions from which one may conclude that a Factor applies.

The negative instances are all the other sentences in the summary. (Brüninghaus & Ashley, 2001).

In current Ph.D. dissertation work, we are testing the hypothesis that by automatically generalizing the training instances to reflect the argument roles of the participants and objects, by schematizing their relationships, and by roughly demarcating the scope of negation, a program can learn to identify known Factors in new texts and facilitate automated indexing. For example, the *Mason* problem above contained the following sentence from which one may conclude that Factor F1, Disclosure-in-Negotiations (d), applies: “Mason disclosed part of the recipe to Randle in exchange, Mason claimed, for a promise that Mason and his band would be used in a sales promotion.” As a training instance, this sentence is likely to be much more effective if one can replace specific names of parties and their products with role-playing concepts like “plaintiff,” “defendant,” and “plaintiff’s product,” and also simplify by extracting patterns, as in, “Plaintiff disclosed part of the recipe to defendant in exchange for a promise that plaintiff and his band would be used in a sales promotion.” Pattern extractions are performed with an adapted version of Ellen Riloff’s Information Extraction (IE) system Autoslog and its Sundance parser (Riloff, 1996). We hope to demonstrate empirically that such generalized training examples can better capture the pattern of concepts associated with a Factor and that the learned decision trees better discriminate positive and negative instances of Factors (Brüninghaus & Ashley, 2001).

5.0 Recommendations for a Civil Law Jurisdiction Contemplating Uses of Precedents.

Are the AI & Law models relevant in civil law jurisdictions? Conceptually, they may be helpful in illustrating some basic ways in which common law attorneys reason with cases: drawing legal inferences from fact-based comparisons of cases, testing hypotheses about who should win a problem against cases, explaining away counterexamples and distinguishing cases, and downplaying or emphasizing distinctions. Of course, the goal of the AI & Law research discussed in the last section has been to model these tasks well

enough to assist common law attorneys in performing them; human attorneys perform these and related tasks in many sophisticated ways that the models cannot yet perform.

In addition, as noted above, there are reasons why civil law jurisdictions may decide to use precedents in some form in legal reasoning, if only to take better advantage of computerized, full-text legal information retrieval. As argued above, a full-text legal information system using Bayesian inference networks could assist practitioners in information retrieval whether the goal is to support retrieving Abstract Precedents or Fact-Based Precedents. Even if the goal is to retrieve cases only for the abstract rules or principles to which they refer, over the long term, courts could improve the effectiveness of information retrieval tools by reporting more fully the facts of the cases in published case opinions. Of course, if the goal is to support retrieval of Fact-Based Precedents, then it is imperative that judges report the facts of a case fully in their opinions.

In either event, implementing a full-text legal information system is an important first step. While it requires an investment in acquiring information retrieval software and in organizing the assembly of the database of texts, it does not require research advances.

The question remains, however, should researchers in a civil law jurisdiction pursue work on implementing AI & Law models of case-based legal reasoning in a civil law context? To the extent that practitioners and judges are interested in retrieving Fact-Based Precedents and in using tools to assist them in drawing legal inferences from a comparison of the facts of problem and cases, the answer may well be, "Yes"!

By now, the tradeoff of benefits and costs in a computational model of legal reasoning is clear. In a legal domain where a Factor-based representation is appropriate, an AI & Law model can support Fact-Based Precedent retrieval for the purposes of automating the drawing of legal inferences from case comparisons, predicting the outcomes of legal problems, and generating the best competing arguments by analogy. On the other hand, the AI & Law models depend on manual entry of cases and problems, and their case databases are

quite small as a result. Techniques for integrating the AI & Law models with full-text legal information systems and tools for extracting information from case texts are under development. The problem of connecting with case texts is especially difficult in the US, because of the length of the legal opinions, the complexity of the prose, and the lack of a standardized structure for legal opinions.

Perhaps, the most interesting opportunity for the judicial and AI & Law communities in a jurisdiction new to case-based legal reasoning is in cooperating in the design of standards for the presentation of factual descriptions and discussions of law in case opinions. Assuming that a civil law jurisdiction's case opinions will begin to include more lengthy descriptions of the cases' facts, it may be possible to invent structures and standards that will assist AI & Law programs to connect more easily to the opinion texts and with full text information retrieval systems. To the extent that civil law judges are just beginning to report fact descriptions, they may be willing to write decisions in a manner that would facilitate automated legal reasoning with the resulting cases. Such standards might include:

- A standard structure for case opinions, with standardized demarcations of the parts of the opinion that contain descriptions of the facts, descriptions of the law, and application of the law to the facts, or to such other structures that make sense in the evolving legal context.
- Standard ways for indicating the parties' roles in the law suit, the claims involved, who won them, and depending on the claim, particular information that most claims of that type involve.
- Standardized ways of describing factors present in a case. Assuming that stereotypical patterns of facts play a role in judicial decision making in particular legal domains, it would be helpful to demarcate, for instance, the factors that favor the plaintiff, those that favor the defendant, and the issues for which they are relevant. Conceivably, it may also be possible substantively to mark-up or tag the text to indicate which facts are critical for which conclusions about which legal issues.

Developing techniques for substantively marking up legal opinions in a way that would support automated summarization and indexing is still an area for research (See, e.g., Grover, et al., 2003). The most important point, however, is to recognize the opportunity presented. Assuming that judicial patterns of opinion writing will change to incorporate greater coverage of case facts, the AI & Law community may be able (1) to determine reasonable standards that make the most sense in light of the evolving use of precedents in judicial reasoning and that facilitate automated extraction of information for indexing and inference, and (2) to help institutionalize conformance to the standards before judicial patterns of opinion writing become settled. This may entail providing judges with networked opinion-writing environments that incorporate the tools and standards into the word-processing infrastructure.

6.0 Conclusions

The relevance of case-based computational models of legal reasoning in the context of a civil law jurisdiction depends on many considerations. As discussed above, evidence suggests that judges in civil law jurisdictions do reason with legal cases, but they reason with cases in a very different way from their common law counterparts. There are, however, some reasons to believe that increasingly civil law judges will find it worthwhile to begin more fully to report the facts of a case in their opinions. For one thing, computerized full-text legal information retrieval, a tool that can benefit civil law judges and practitioners as much as anyone, works better with fuller fact descriptions, even if the goal is to retrieve only the principles and abstract rules a court relies upon. For another, international treaties and other considerations suggest that civil law judges increasingly will need to compare current problems to past decided cases for purposes of drawing legal inferences. To the extent this is true, case-based computational models of legal reasoning offer techniques for improving upon the ability of full-text legal information systems to process retrieved cases in an intelligent way that reflects their significance in legal arguments.

Finally, to the extent that judges in a civil law jurisdiction have not yet adopted standards for reporting the facts of cases, there may be an important

opportunity for AI & Law researchers to help determine those standards with an eye toward helping their computational models process the cases intelligently and automatically.

7.0 References

- Aleven, V. (2003) Using Background Knowledge in Case-Based Legal Reasoning: A Computational Model and an Intelligent Learning Environment. *Artificial Intelligence*. Special Issue on Artificial Intelligence and Law. In Press.
- Aleven, V. (1997) *Teaching Case-Based Argumentation Through a Model and Examples* (Ph.D. dissertation, University of Pittsburgh Graduate Program in Intelligent Systems) (unnumbered technical report, Learning Research and Development Center).
- Aleven, V. and Ashley, K. (1997) "Teaching Case-Based Argumentation through a Model and Examples: Empirical Evaluation of an Intelligent Learning Environment", in *Artificial Intelligence in Education, Proceedings AI-Ed 97 World Conference on Artificial Intelligence in Education*, pp. 87-94. (August. B. du Boulay and R. Mizoguchi, ed. IOS Press).
- Aleven, V. and Ashley, K. D. (1994) "An Instructional Environment for Practicing Argumentation Skills". In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*. pp. 485-492. Seattle, WA. July. AAAI Press: Menlo Park, CA.
- Ashley, K. (1990) *Modeling Legal Argument: Reasoning with Cases and Hypotheticals* (The MIT Press).
- Ashley, K.D. (1987) *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. (1987) Ph.D. Dissertation. COINS Technical Report No. 88 01, Department of Computer and Information Science, University of Massachusetts, Amherst.
- Bench-Capon, T. and Sartor, G. (2001) "Theory Based Explanation of Case Law Domains", in *Proceedings, 8th Int'l Conference on Artificial Intelligence and Law*. pp. 12-21. (Association for Computing Machinery).
- Branting, L.K. (1991) "Building Explanations from Rules and Structured Cases", *Int'l Journal of Man-Machine Studies*, v. 34, pp. 797-837.
- Branting, L.K. (1999) *Reasoning with Rules and Precedents - A Computational Model of Legal Analysis*. Kluwer Academic Publishers.
- Brüninghaus, S. and Ashley, K.D. (2003) Predicting the Outcome of Case-Based Legal Arguments. In: *Proceedings of the Ninth International Conference on Artificial Intelligence and Law (ICAIL-03)*, 234-242.
- Brüninghaus, S. and Ashley, K.D. (2001). "Improving the Representation Of Legal Case Texts With Information Extraction Methods". In the *Proceedings of the Eighth International Conference on Artificial Intelligence & Law (ICAIL-01)*. pp. 42-51. ACM Press.
- Callan, J.P., Croft, W.B. and Harding, S.M. (1992) The INQUERY Retrieval System. In Tjoa, A.M. and Ramos, I. eds., *Database and Expert Systems Applications: Proceedings fo the International Conference in Valencia, Spain*. Pp. 78-83. Springer Verlag. NY.
- Curran, V. G. 2001. Romantic Common Law, Enlightened Civil Law: Legal Uniformity and the Homogenization of the European Union. *7 Columbia Journal of European Law* Winter 63.
- Grover, C., Hachey, B., Hughson, I., Korycinski, C. (2003) Automatic Summarisation of Legal Documents. In: *Proceedings of the Ninth International Conference on Artificial Intelligence and Law (ICAIL-03)*, 243-251.
- Lundmark, T. 1998. Book Review, Interpreting Precedents: A Comparative Study. MacCormick, D.N. & Summers, R.S. (ed.) 46 *American Journal of Comparative Law* 211.
- MacCormick, D.N. & Summers, R. S. (ed.) 1997. *Interpreting Precedents: A Comparative Study* Dartmouth Publishing: Aldershot, U.K.
- Prakken, H. and Sartor, G. (1997) "Reasoning with Precedents in a Dialogue Game", in *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*, pp. 1-9 (Association for Computing Machinery).

- Riloff, E. (1996) Automatically Generating Extraction Patterns from Untagged Text, *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pp. 1044-1049, AAAI Press / The MIT Press. Menlo Park, CA.
- Rissland, E.L. & Daniels, J.J. (1996) "The Synergistic Application of CBR to IR." *Artificial Intelligence Review: Special Issue on the use of AI in Information Retrieval*, 10, pp. 441-475.
- Rissland, E.L. and Skalak, D.B. (1991) CABARET: Statutory Interpretation in a Hybrid Architecture. *International Journal of Man-Machine Studies* 34(6):839-887.
- Rissland, E.L., Skalak, D.B. et al. (1996) "BankXX: Supporting Legal Arguments Through Heuristic Retrieval", *Artificial Intelligence and Law*, v. 4, p. 19.
- Turtle, H.R. (1995) "Text Retrieval in the Legal World" in *Artificial Intelligence and Law*, 3: 5-54. Kluwer: Dordrecht, The Netherlands.
- Zeleznikow, J. Stranieri, A. et al. (1995-1996) "Project Report: Split-Up – A Legal Expert System which Determines Property Division Upon Divorce", *Artificial Intelligence and Law*, v. 3, pp. 267-275.