Legal Reasoning: Computational Models

Henry Prakken
Department of Information and Computing Sciences, Utrecht University
P.O.Box 80089, 3508 TB Utrecht, The Netherlands
and
Faculty of Law, University of Groningen
P.O. Box 716, 9700 AS Groningen, The Netherlands
Email: H.Prakken@uu.nl

Abstract

This article reviews computational models of legal reasoning as they are being developed in the field of artificial intelligence (AI) and law. A theoretical aim of these models is to understand legal reasoning by simulating it in a computer program, while a practical aim is to study how advanced information technology can aid legal practice. First logical deduction is discussed as a necessary but insufficient component of realistic models of legal reasoning. Then models of defeasible legal reasoning and legal argumentation are discussed, which focus on the generation and comparison of reasons or arguments for and against legal claims. Computational models of legal interpretation emphasize the interplay between rules and cases and the role of principles, purposes and values. Computational models of legal proof account for the uncertainty in legal proof in three alternative ways: with Bayesian probability theory, with argumentation and with narrative. Finally, procedural models of legal reasoning are based on the idea that the quality of a legal decision partly depends on how it was reached.

Introduction

This article[1] reviews computational models of legal reasoning as they are being developed in the field of artificial intelligence (AI) and law. Computational models of legal reasoning have been developed for two main purposes. A theoretical aim is to understand legal reasoning by simulating it in a computer program, while a practical aim is to study how advanced information technology can aid legal practice. Both types of research have had considerable success. Computational models of legal argument have advanced the understanding of legal reasoning, while legal knowledge-based systems are in routine use in areas of legal administrative decision making in several countries.

Some may fear that using computational methods in the law enforces mechanical law application, without regard for social consequences or for principles

[1] Some parts of the text of this article are adapted from Prakken and Sartor (2002) or Prakken and Sartor (2014).

of fairness and justice. However, this is a misconception. Computational approaches require clarity of meaning and explicitness of structure, but this should not be confused with mechanical law application or with the view that the only valid source of the law is legal rules. Although some early systems were open to this criticism, there is nothing in a computational approach that requires judges to mechanically apply rules without looking at the consequences of applying the rule. On the contrary, computational models exist of arguing why an applicable rule should not be applied to a case or why a non-applicable rule should be broadened to make it apply. Likewise, there is nothing in a computational approach that insists on the uniqueness or completeness of interpretations of the law. On the contrary, computational models exist to argue for alternative interpretations and about their relative merits. Finally, computational models of legal reasoning exist that acknowledge cases, principles or opinions in society about justice and fairness as valid sources of the law.

Models of legal reasoning must be diverse. A legal case has various aspects, each with its own modes of reasoning: determining the facts, classifying the facts under legal concepts or conditions, and deriving legal consequences from the thus classified facts. When determining the facts, the modes of reasoning are often probabilistic and may involve reasoning about causation and about mental attitudes such as intent. Classifying the facts under legal concepts involves interpretation. Here the prevailing modes of reasoning are analogy, appeals to precedent or policy, and the balancing of interests. Finally, when deriving legal consequences from the classified facts, the main modes of reasoning are deductive but leave room for exceptions to rules, either statutory or based on principle and purpose, and have room for conflict based on the general hierarchy of legal systems, where rules derive from different sources. Models of legal reasoning also depend on the task that is modeled, for example, advice, advocacy or decision making. This entry will primarily focus on the context of a court proceeding, where a judge has to adjudicate a conflict between two adversaries.

About any computational model of reasoning three questions can be asked:

1. What is the underlying theory of reasoning?

2. Can the knowledge required by the reasoning be represented in the computer?

3. Can the computer reason efficiently with this knowledge?

This article will primarily focus on the first question but the other two will not be totally ignored. The third question is addressed in other areas of AI, and the resulting techniques can be used in AI and law when available. The problem addressed by the second question has so far been the bottleneck, hindering scaling computational models of legal reasoning up to practical use. Where it has been solved, (either by explicit representation or by automatically learning the knowledge from data), feasible and successful applications have resulted.

Legal reasoning as deduction

A simple but naive model of legal reasoning sees it as logically deducing legal consequences from a precisely stated body of facts and legal rules. On this account, once a legal text and a body of facts have been clearly represented in a logical language, the valid inferences are determined by the meaning of the representations and so techniques of automated deduction apply. However, this mechanical approach leaves out most of what is important in legal reasoning. To start with, the law is not just a conceptual or axiomatic system but has social objectives and social effects, which must be taken into account when applying the law. Moreover, legislators can never fully predict in which circumstances the law has to be applied, so legislation has to be formulated in general and abstract terms, such as 'duty of care', 'misuse of trade secrets' or 'intent', and qualified with general exception categories, such as 'self defence', 'force majeure' or 'unreasonable'. Such concepts and exceptions must be interpreted in concrete cases, a process which creates room for doubt and disagreement. This is reinforced by the fact that legal cases often involve conflicting interests of opposing parties. The prosecution in a criminal case wants the accused convicted while the accused wants to be acquitted. The plaintiff in a civil law suit wants to be awarded compensation for damages, while the defendant wants to avoid having to pay. The tax authority in a tax case wants to receive as much tax as possible, while the taxpayer wants to pay as little as possible. Both aspects of the law, i.e., the tension between the general terms of the law and the particulars of a case, and the adversarial nature of legal procedures, cause legal reasoning to go beyond the meaning of the legal rules. It involves appeals to precedent, principle, policy and purpose, as well as the consideration of reasons for and against drawing conclusions. For these reasons law relies not just on deduction but on argument.

Nevertheless, deduction is still an essential ingredient in legal reasoning. To start with, the principle of legal certainty requires that the meaning of the law is not arbitrarily set aside. And when arguing that the law should be set aside in a certain case because applying it would have undesirable effects, the effects of applying the law must be deduced. Also, contemplating alternative interpretations of a legal concept involves deducing how the various interpretations affect the outcome of a case. Moreover, deductive techniques have been practically successful, especially in the application of knowledge-based systems in large-scale processing of administrative law, such as social benefit law. Such systems apply computational representations of legislation to the facts as interpreted by the human user. The use of such systems has been proved to greatly reduce two major sources of errors in the processing of social benefit applications by 'street-level bureaucrats': their incomplete knowledge of the relevant regulations and their inability to handle the often complex conditions of the regulations (Svensson; 2002). These systems leave it to the user (the official with the authority to make a decision) to decide whether to accept the system's recommendation or to deviate from it on non-statutory grounds. Thus these systems do not enforce mechanical application of the law, at least as long as the correctness of their recommendations is not taken for granted. We can conclude that, even though legal reasoning does not equate with logical deduction, the study of the deductive aspects of legal reasoning is important for both theoretical and practical reasons.

Much research on these aspects exist and we can discuss only some of it here. A seminal study of the logical structure of legislation was Sergot et al. (1986)'s modeling of the British Nationality Act as a logic program, which showed how logic programming (and logic more generally) enables intuitively appealing representations that can be directly deployed to generate automatic inferences.

Here is how a legislative rule from the British Nationality act is modeled as a logical clause:

```
x is a British citizen
IF x was born in the U.K. AND x was born on date y
AND y is after or on commencement
AND x has a parent who qualifies under 1.1 on date y
```

Sergot et al. also showed how further clauses can be represented that determine the conditions under which the conditions in the legal rules hold. This work has greatly influenced subsequent research on knowledge-based systems for processing legislation.

Other important work is on the logical characterization of the various legal positions or relations that may exist between normative agents as regards an action or state of affairs, in terms of their rights, obligations and permissions. This work enables the expression of subtle but important distinctions, such as whether a patient's normative position with regard to access to his medical data, besides the patient's permission to access the data, does or does not include the medical staff being forbidden to impede the patient's access. In the first case, the patient has the right to access his medical data; in the second case he is merely permitted to do so. Sergot (2001) presents a computational method for exhaustively specifying and computing the alternative normative positions of two parties concerning an action. Such methods can be useful, for example, for drafting or disambiguating security or access policies.

Among other work on deductive aspects of legal reasoning we mention methods for representing systems of legal concepts (Bench-Capon and Coenen; 1992) and methods for representing change in legal systems, normative positions or legal relations over time (Rotolo and Governatori; 2009).

Exceptions and rule hierarchies

Any attempt to computationally model legal regulations soon meets two very common structural features of legal regulations that provide obstacles for deductive techniques: the separation of general rules and exceptions, and the use of hierarchies over legislative sources to resolve conflicts between different regulations. With exceptions the problem is that the law often makes legal effects dependent on the assumption that they do not occur as long as there is no evidence for their occurrence. For example, Section 3:32 of the Dutch civil code states that persons have the capacity to perform legal acts (this means, for instance, that they can engage in contracts or sell their property), unless the law provides otherwise. One place in which the law provides otherwise is Section 1:234 of the Dutch civil code, which states that minors have the capacity to perform legal acts if and only if they have consent from their legal representative. If we take a person named John and then consider these two rules, it follows that John has the capacity to perform legal acts. However, this inference is presumptive or defeasible in that it can be invalidated by further information. For example, if evidence is provided that John is a minor, then the conclusion that he can perform legal acts is invalidated unless evidence is also provided that John has consent of his legal representative (usually one of his parents).

The making and retracting of such defeasible inferences goes beyond standard deductive models of reasoning. One way in which AI and law has tackled this problem is by the use of so-called non-monotonic logics. Such logics, which are a result of general AI research, allow the making of inferences subject to evidence to the contrary, and the retraction of such inferences when evidence to the contrary comes in. These logics can also model reasoning with hierarchies. Models of deduction have no room for conflict resolution or hierarchies: when a body of information is logically inconsistent, then standard deductive methods simply come to a halt. In contrast, nonmonotonic logics have been designed in which inconsistencies can be resolved by defining preference relations on rules, and these preferences can be used to express legal hierarchies. This too makes the reasoning defeasible, since new information may make a conflicting rule of higher rank applicable that overrules a previously applicable rule with lower rank. Nonmonotonic logics have been shown useful in modeling legislative rule-exception structures and legislative hierarchies (Prakken and Sartor; 1996), and in modeling legal presumptions and notions of burdens of proof (Prakken and Sartor; 2009; Gordon and Walton; 2009).

Although nonmonotonic techniques technically deviate from deductive logic, their spirit is still the same, namely, of deriving consequences from clear and unambiguous representations of legal rules and facts. The logical consequences of a representation are still clearly defined. While technically most nonmonotonic logics allow for alternative conclusion sets (which capture cases in which conflicts cannot be resolved), in legal practice statutory rule-exception structures and legislative hierarchies still usually yield unambiguous outcomes. More often, conflicts arise not from competing norms but from the variety of ways in which they can be interpreted. A real challenge for deductive accounts of legal reasoning is the gap between the general legal language and the particulars of a case. Because of this gap, disagreement can arise, and it will arise because of the conflicts of interests between the parties.

Rule-based accounts of legal interpretation

At first sight, it might be thought that disagreements about interpretation are resolved in concrete cases by courts, so that additional interpretation rules can be found in case law. If different courts disagree on an interpretation, such disagreements could be represented with nonmonotonic techniques for handling conflicting rules. This was indeed the approach taken by Gardner (1987), who designed a program for so-called "issue spotting" in law school contract law exam problems. Given an input case, the task of the program was to determine which legal questions involved were easy and which were hard, and to solve the easy ones. If all the questions were found easy, the program reported the case as clear, otherwise as hard. The system contained domain knowledge of three different types: legal rules, common-sense rules (e.g. on the interpretation of utterances like "Will you supply ...?"), and rules extracted from cases. The program considered a question as hard if either "the rules run out", or different rules or cases point at different solutions, without there being any reason to prefer one over the other. Before a case was reported as hard, conflicting alternatives were compared to check whether one is preferred over the other. For example, case law sets aside conflicting legal rules or common-sense interpretations of legal concepts. Gardner's program was reconstructed by Gordon (1991) with techniques from nonmonotonic logic. Another approach is Hage & Verheij's Reason-Based Logic (Verheij et al.; 1998), which models reasoning as the weighing of reasons for and against a conclusion, where rules are just one source of reasons, along with, for example, precedents or principles. Thus they could, like Gardner, model a statutory rule being set aside not by a conflicting statutory rule but by other considerations.

A problem with rule-based approaches is that rules derived from precedents are often specific to the case, so a new case will rarely exactly match the precedent, and techniques for handling conflicting rules fall short. Instead, reasoning forms are called for in which case-law rules can be refined and modified. Here factors and reasons play an important role, and analogies between cases are drawn and criticized.

Interpreting legal concepts with cases

What are the reasoning forms "when the rules run out"? They can be found in their most explicit form in common-law jurisdictions, in which judicial precedents can be legally binding beyond the decided case, so that court decisions legally constrain the decision in new cases. This leads to reasoning forms where legal rules are formulated by courts in the context of particular cases and are constantly refined and modified to fit new circumstances that were not taken into account in earlier decisions. These reasoning forms can to a lesser extent also be found in civil-law jurisdictions, since interpretations of the law by higher civil-law courts, even though strictly speaking not binding beyond the decided case, still tend to be followed by lower courts.

Perhaps the earliest work in AI and law on legal interpretation was the TAXMAN II project of McCarty (1977, 1995)). According to McCarty (1995, p. 285), "The task for a lawyer or a judge in a "hard case" is to construct a theory of the disputed rules that produces the desired legal result, and then to persuade the relevant audience that this theory is preferable to any theories offered by an opponent". In McCarty's approach, legal concepts have three components: firstly, a (possibly empty) set of necessary conditions for the concept's applicability; secondly, a set of instances ("exemplars") of the concept; and finally, a set of rules for transforming one case into another, particularly for relating "prototypical" exemplars to "deformations". According to McCarty, the way lawyers typically argue about application of a concept to a new case is by finding a plausible sequence of transformations which maps a prototype, possibly via other cases, onto the new case. McCarty's particular modeling of this approach has not found many followers but his general claim that legal interpretation is theory construction has been very influential.

Much AI & law work on the interpretation of legal concepts centers around the notion of a factor, an idea going back to the HYPO system of Ashley (1990) and the CATO system of Aleven (2003). Factors are abstractions of fact patterns that favor (pro factors) or oppose (con factors) a conclusion. Factors are thus in an intermediate position between the specific facts of a case and the legal predicates to which such facts may be relevant. For example, in CATO, which like HYPO argues about misuse of trade secrets, some factors pro misuse are that a non-disclosure agreement was signed, that the plaintiff had made efforts to maintain secrecy and that the copied product was unique; and some factors con misuse are that disclosures were made by the plaintiff in negotiations and that the information was reverse-engineerable.

The HYPO and CATO systems are meant to model how lawyers make use of past decisions when arguing a case. They do not compute an 'outcome' or 'winner' of a dispute; instead they are meant to generate debates as they could take place between 'good' lawyers. HYPO generates disputes between a plaintiff and a defendant of a legal claim concerning misuse of a trade secret. Each move conforms to certain rules for analogizing and distinguishing precedents. These rules determine for each side which are the best cases to cite initially, or in response to the counterparty's move, and how the counterparty's cases can be distinguished. According to the theory described by Ashley (1989) and implemented in CATO a case is represented as a set of factors for a decision and a set of factors against that decision, plus the decision that resolves the conflict between the competing factors. A case is citable for a side if it has the decision wished by that side and shares with the Current Fact Situation (CFS) at least one factor which favors that decision. Thus citable cases do not have to exactly match the CFS, which is a way of coping with the case-specific nature of case law decisions. A citation can be countered by a counterexample, that is, by producing a citable case that has the opposite outcome. A citation may also be countered by distinguishing, that is, by indicating a factor in the CFS that is absent in the cited precedent and that supports the opposite outcome, or a factor in the precedent that is missing in the CFS and that supports the outcome of the cited case. HYPO also allows for multi-valued factors, called dimensions, to vary the degree to which a factor promotes a certain outcome. For example, a dimension is the number of people to which a trade secret has been disclosed, or the extent to which security measures were taken. A boolean factor is then a specific value of a dimension. Dimensions allow an additional way to distinguish a precedent, namely, on a shared pro-decision factor that more strongly favours the decision in the precedent than in the CFS.

The HYPO system was embedded in the CABARET system of Skalak and Rissland (1992), which modeled the use of precedents to confirm or contest the application of a rule. The system implemented a hierarchical argumentation model in terms of strategies, moves and primitives. A strategy is a broad characterization of how one should argue, given one's particular viewpoint and dialectical situation. A move is a way to carry out the strategy, while a primitive is a way to implement a move. For example, when one wants to apply a rule, and not all of the rule's conditions are satisfied, then a possible strategy is to broaden the rule. This strategy can be implemented with a move that argues, using an analogized precedent, that the missing condition is not really necessary. This move can in turn be implemented with HYPO's ways to analogize cases. Similarly, CABARET also permits arguments that a rule which *prima facie* appears to cover the case should not be applied to it. Here the strategy is discrediting a rule and the move may consist of analogizing a case in which the rule's conditions were met but the rule was not applied. Again the move can be implemented with HYPO's ways to analogize cases.

Aleven (2003) developed HYPO (without dimensions) into the CATO system for teaching case-based argumentation skills to law students, also in the trade secrets domain. CATO's main new component was a 'factor hierarchy', which expresses expert knowledge about the relations between the various factors:

more concrete factors are classified according to whether they are a reason for or against the more abstract factors to which they are linked; links are given a strength (weak or strong), which can be used to solve certain conflicts. Thus the factor hierarchy can be used to explain why a certain decision was taken, which in turn facilitates debate on the relevance of differences between cases.

For instance, the hierarchy positively links the factor *Security measures taken* to the more abstract concept *Efforts to maintain secrecy*. Now if a precedent contains the first factor but the CFS lacks it, then not only can a citation of the precedent be distinguished on the absence of *Security measures taken*, but also this distinction can be emphasized by saying that thus no efforts were made to maintain secrecy. However, if the CFS also contains a factor *Agreed not to disclose information*, then the factor hierarchy enables downplaying this distinction, since it also positively links this factor to *Efforts to maintain secrecy*: the party that cited the precedent can say that in the current case, just as in the precedent, efforts were made to maintain secrecy.

Perhaps the most elaborate representation of cases was produced in Branting (1999)'s GREBE system in the domain of industrial injury, where cases were represented as semantic networks. The program matched portions of the network for the new case with parts of the networks of precedents, to identify appropriate analogies.

Various logical accounts of factor-based reasoning with cases have been given. Here the nonmonotonic logics that were originally used to model reasoning with exceptions and legal hierarchies turned out to be useful for this task also. Most of this work applies argumentation-based non-monotonic logic, in which rules can be repeatedly applied to construct arguments, and in which conclusions are drawn by comparing alternative arguments for or against conclusions. (But some accounts are formulated in terms of weighing sets of reasons for or against a conclusion.) The key idea in these approaches is that existing case decisions give rise to conflicting rules (or conflicting sets of reasons) plus a preference expressing how the court resolved this conflict. In the notation of Prakken and Sartor (1998):

r_1: *Pro-factors* $\Rightarrow$ *Decision*
r_2: *Con-factors* $\Rightarrow$ *Not D*ecision
$r_1 > r_2$

The priority expresses the court's decision that the pro factors in the body of rule r_1 together outweigh the con factors in the body of rule r_2. This approach allows for 'a fortiori' reasoning in that adding factors to a pro-decision rule or removing factors from a con-decision rule does not affect the rule priority. It allows for analogical uses of a pro-decision rule by deleting pro factors, which is called broadening a rule. A broadened rule does not inherit the priority relations of the rule from which it is obtained. So if rule r_1 in our schematic example is broadened to rule $r_{1'}$ by deleting one of the pro-factors in rule r_1, then one cannot conclude from the priority $r_1 > r_2$ in the precedent that $r_{1'} > r_2$, since the deleted factor might have been essential in preferring r_1 over r_2. So a decision argued for with a broadened rule is not (defeasibly) implied by a body of precedents.

An interesting question is under what conditions such a decision is still allowed by the body of precedents. This question is addressed by Horty (2011). Consider the *Keeble* case from the common law of property, part of a well-known series of cases on ownership of wild animals that are being chased. In *Keeble* a pond owner placed a duck decoy in his pond with the intention to sell the caught ducks for a living. Defendant used a gun to scare away the ducks, for no other reason than to damage plaintiff's business. Here the court held for plaintiff. The pro-plaintiff factors were that plaintiff was hunting for a living (*PlLiving*) and was hunting on his own land (*OwnLand*). The single pro-defendant factor was that the animals were not yet caught (*N*ot Caught). The issue is whether plaintiff became the owner of the ducks (*Owner*). In the notation of Prakken and Sartor (1998) we have:

Keeble
k_1: *PlLiving*, *OwnLand* $\Rightarrow$ *Owner*
k_2: *Not Caught* $\Rightarrow$ *Not Owner*
$k_1 > k_2$

In another precedent, *Young*, both plaintiff and defendant were fishermen fishing in the open sea. Just before plaintiff closed his net, defendant came in and caught the fishes with his own net. Here not only the plaintiff but also the defendant was hunting for a living (*DefLiving*). Then we have

Young
y_1: *PlLiving* $\Rightarrow$ *Owner*
y_2: *Not Caught*, *DefLiving* $\Rightarrow$ *Not Owner*

To decide *Young* for plaintiff, the required priority $y_1 > y_2$ cannot be based on the precedent, since y_1 lacks one antecedent of k_1, and also since y_2 adds a con factor to k_2. However, deciding *Young* in accordance with *Keeble* is still allowed by *Keeble*, since that leaves the decision in the precedent unaffected. Horty calls this *following* the precedent.

This is different if the case base also contains a second precedent, which is almost like *Keeble* except that the defendant also hunted for a living and in which the defendant won.

Precedent 2
p_1: *PlLiving*, *OwnLand* $\Rightarrow$ *Owner*
p_2: *Not Caught*, *D*efLiving $\Rightarrow$ *Not Owner*
$p_2 > p_1$

The priority $p_2 > p_1$ then implies $y_2 > y_1$, since y_2 has the same con-decision factors as p_2 while y_1 lacks one pro-decision factor of p_1. But $y_2 > y_1$ is inconsistent with $y_1 > y_2$, so deciding *Young* as *Keeble* is not allowed by the extended case base, since it would amount to *overruling* the second precedent.

In Horty's approach not all deviations from a precedent are overrulings. Suppose that precedent 2 is not in the case base but is a new case, so its decision is not yet known. Then $p_2 > p_1$ is consistent with $k_1 > k_2$, so deciding the new case differently than in the precedent is allowed by the case base, since it leaves all decisions in precedents unaffected. Horty here says that deciding con the original decision in the new case *distinguishes* the precedent. Horty has thus given precise logical formalizations of the important common-law notions of following, distinguishing and overruling a precedent.

These approaches have been extended by allowing for chains of reasoning from more concrete to more abstract factors, so that detailed reasoning about how the factors relate to a conclusion can be modeled (cf. CATO's factor hierarchy). For example, Prakken and Sartor (1998), using an idea of Loui and Norman (1995), represent cases in general as multi-step arguments for and against conclusions, where rule priorities are expressed for each pair of conflicting rules, and where any rule or priority of a precedent can be used in a new case. This makes it possible to argue in various ways that in a new case some of the arguments pro the decision cannot be constructed or new arguments con the decision can be constructed, so that the new case must have a different outcome (cf. Loui and Norman (1995)).

The work on chains of reasoning and dialectical structures does not explain why the most abstract factors are factors pro or con a decision. Berman and Hafner (1993) argued that often a factor can be said to favor a decision by virtue of the purposes served or values promoted by taking that decision because of the factor. A choice in case of conflicting factors is then explained in terms of a preference ordering on the purposes, or values, promoted or demoted by the decisions suggested by the factors. Cases can then be compared in terms of the values at stake rather than on the factors they contain.

Bench-Capon (2002) computationally modeled this approach and illustrated it with the series of cases from American property law on ownership of wild animals that are being chased. For example, plaintiff in *Keeble* could argue that people should be protected when pursuing their livelihood, since society benefits from their activities. He could also argue that he was hunting on his own land, so that the value of protection of property is another reason why he should win. Defendant in *Keeble* could argue that since plaintiff had not yet caught the ducks, he had no right to the ducks, since if such rights depended on who first saw the animals, there would be no clear criterion and the courts would be flooded with cases. Thus defendant argues that deciding for him promotes the value of legal certainty. Since plaintiff won in *Keeble*, we can on this interpretation of the case say that the court found that the combination of the values of protecting property and protecting the pursuit of livelihood outweighs the single value of legal certainty. This value preference can be cited in new cases where the same values are at stake (but note that in *Young* the value of protecting the pursuit of livelihood is cancelled out by the fact that in *Young* the defendant was also pursuing his livelihood).

A limitation of this approach is that often the extent to which a value is promoted or demoted is also important. Even if a new case contains the same values as a precedent, the degrees of promotion or demotion may be different, which can give rise to arguments against the citation. This could explain why in legal practice almost no generally valid hierarchies of values are accepted: the point is that different degrees of promotion or demotion give rise to different value preferences in different cases.

Bench-Capon and Sartor (2003) embedded Bench-Capon (2002)'s approach in a method for constructing theories that explain a given set of cases. They modeled theory construction as an adversarial process, where both sides take turns to modify the theory so that it explains the current case in the way they want. The process starts with a set of statements about which factors promote which values and a set of cases represented in terms of factors and an outcome. Then the theory is constructed by creating rules plus rule preferences derived from value preferences. Chorley and Bench-Capon (2005) study various ways to automate this process.

The attention paid to the role of value and purpose led to accounts of legal interpretation as a decision problem, namely, as a choice between alternative interpretations on the basis of the likely consequences of these interpretations in terms of promoting and demoting values. Thus current AI models of argumentation for decision making can be applied. This approach in fact regards legal reasoning as a form of what philosophers call practical reasoning. It compares arguments about what decision to take in terms of the various positive and negative effects of the decision alternatives on goals and values, in terms of the relative preferences between these goals or values, and sometimes also in terms of the extent to which they are promoted or demoted. This line of research has especially been pursued by Bench-Capon and Atkinson and colleagues, e.g. in Atkinson et al. (2005), applying it both to legal interpretation and to legislative and policy debates.

A practical-reasoning account of legal interpretation can be applied to model the problem of a court deciding whether to follow or distinguish a precedent when a new factor arises (assuming that distinguishing is consistent with the body of precedents). For example, deciding *Young* for the defendant can be reconstructed as follows. First two practical-reasoning arguments for following or distinguishing *Keeble* are stated. Then the argument for distinguishing *Keeble* is preferred over the argument for following *Keeble* on the ground that following *Keeble* only promotes the value of protecting livelihood while distinguishing *Keeble* in addition promotes the value of legal certainty.

The work reviewed in this section does not amount to a full theory of interpretative legal argument. One thing hardly addressed in the above models is the step from facts to factors. Further study is needed of how legal reasoning actually bridges the gap between the concrete facts and the more abstract factors in a case.

Establishing the facts of a case

While legal education and scholarship mostly focuses on reasoning with and about the law, in practice most cases are decided on the facts, so insight as to how facts can be proven is crucial for legal practice. AI & law research has addressed two main questions: which model of rational proof can best be applied to the law, and what is the logical nature of important legal evidential constructs like burdens of proof and presumptions? We address these issues in turn.

Models of legal proof

Theoretical models of rational legal proof are generally of three kinds: probabilistic, story-based, or argument-based. All three approaches acknowledge that evidence cannot provide watertight support for a factual claim but always leaves room for doubt and uncertainty, but they account for this in different ways. Probabilistic approaches express uncertainty in terms of numerical probabilities attached to hypotheses given the evidence. Often a Bayesian approach is taken, nowadays more and more with Bayesian networks (Fenton and Neil; 2011). Probabilistic approaches are by no means uncontroversial. One objection is that in legal cases the required numbers are usually not available, either because there are no reliable statistics, or because experts or judges are unable or reluctant to provide estimates of probabilities. Another objection is that probability theory imposes a standard of rationality that cannot be attained in practice, so that its application would lead to more instead of fewer errors. To overcome these and other claimed limitations of probabilistic models, argumentation-based and story-based models have been proposed.

Story-based approaches go back to the work of the psychologists Bennett and Feldman (1981), who observed that the way judges and prosecutors tend to make factual judgments is not by probabilistic or logical reasoning but by constructing and comparing alternative plausible stories about what might have happened. In these approaches the story that best explains the evidence must, if it does so to a sufficient degree, be adopted as true. Room for doubt is provided since some as yet unknown story might be the true one or since new evidence might make another of the known stories the best one. Thagard (2004) sketches how this approach can be computationally modeled as inference to the best explanation.

Both Bayesian and story-based approaches reason from hypotheses to the evidence in that, to assess alternative hypotheses, they model how likely the evidence is under the various hypotheses. In contrast, argumentation-based approaches reason from the evidence to the hypothesis, by stepwise building evidential arguments from the available evidence to the hypotheses. Room for doubt is provided since additional evidence might give rise to new arguments that defeat the currently undefeated arguments. These approaches go back to Wigmore (1931)'s charting method of legal evidence, with which alternative arguments from evidence to hypotheses can be graphically displayed and thus sources of doubt in these arguments can be revealed. An important source of doubt is the empirical generalizations needed to justify the various steps in an evidential argument, such as that witnesses usually speak the truth. Generally, three sources of doubt in witness testimonies are considered: the witness has reason to lie, the witness has flawed memory, or the witness could not reliably observe, hear, feel or smell what he testified to. Bayesian and story-based approaches account for this by considering four hypotheses: the witness statement is correct, the witness lied, the witness misremembered, or the witness observed incorrectly. They would then estimate how likely the fact that the witness made his statement is under these alternative hypotheses. By contrast, argumentation approaches would defeasibly infer from the witness statement that the accused did what the witness testified to, and they would account for sources of doubt as giving rise to counterarguments.

Bex et al. (2003) model the Wigmorean approach in terms of logical frameworks for argumentation. Evidential reasoning proceeds by applying various argument schemes for evidential reasoning to the evidence, such as schemes for perception, memory, induction, applying generalizations, reasoning with testimonies, and temporal persistence. The so-called critical questions of these schemes specify exceptions to these reasoning forms.

Argumentation approaches are good for modeling how evidence can be related to hypotheses and for revealing sources of doubt in evidential arguments, but less good for providing a clear overview of masses of evidence. This is better done in the story-based approach, which formulates scenarios with a central action made plausible by the context. To combine the strengths of both approaches, Bex (2011) proposed a hybrid theory of legal evidential reasoning. The idea is that stories or scenarios are connected events while evidence for these events and support for their connections is provided by argumentation with argument schemes. Then Bex formulated various criteria for the internal plausibility and coherence of a story and for how well it is supported by the evidence.

Burdens of proof and presumptions

Legal proof is not a scientific problem but takes place under legal and practical constraints. The parties involved in a proceeding have limited technical and financial resources, while a decision has to be reached within reasonable time and in a fair way. Legal systems have developed ways to deal with these constraints, such as presumptions and burdens of proof. Research has been done on modeling these notions with techniques from nonmonotonic logic and argumentation (Prakken and Sartor; 2009; Gordon and Walton; 2009). Briefly, the idea is that a burden of proof for a claim is fulfilled if at the end of a proceeding the claim is acceptable according to the argumentation logic applied to the then available evidence. Some care has to be taken here to model that in the law burden of proofs for different sub-issues (including both burdens of production and burdens of persuasion) can be distributed over the parties. Another challenge is to account for the fact that different kinds of legal issues can have different standards of proof. For example, in common-law jurisdictions claims must in criminal cases be proven 'beyond reasonable doubt' while in civil cases usually proof 'on the balance of probabilities' suffices. In logical accounts defeat between arguments is an all-or-nothing matter, but proponents of probabilistic approaches might say that this fails to respect that defeat between evidential arguments is a matter of degree. Gordon and Walton (2009) and Prakken and Sartor (2011) study ways to reconcile the gradual nature of proof standards with the all-or-nothing nature of defeat between arguments in argumentation logics.

Legal presumptions obligate a fact finder to draw a particular inference from a proved fact. Typical examples are a presumption that the one who possesses an object in good faith is the owner of the object, or a presumption that when a pedestrian or cyclist is injured in a collision with a car, the accident was the driver's fault. Some presumptions are rebuttable while others are irrebuttable. Prakken and Sartor (2009) argue that rebuttable presumptions can be logically interpreted as defeasible rules, so that reasoning with rebuttable presumptions can be modeled with a nonmonotonic logic.

Legal reasoning and legal procedure

Legal reasoning usually takes place in the context of a dispute between adversaries, within a prescribed legal procedure. This makes the setting inherently dynamic and multi-party. The facts and theories are not given at the start of a case, but the adversaries advance their points of view and provide their evidence at various stages, and they accept, reject or challenge their opponent's claims when these are made. The adjudicator can allocate burdens of proof and rule about admissibility of evidence or arguments during a proceeding, before s/he decides the dispute in the end. All these things mean that the quality of a legal decision not only depends on its grounds but also on how it was reached. It is therefore relevant to study how legal procedures contribute to the quality of a decision. In short, legal procedures should promote rational, fair and effective dispute resolution.

AI & Law research has studied these issues by embedding argumentation models like the ones discussed above in dialogue models of legal procedures. The procedures modeled are not actual but artificial ones, often simplified compared to actual legal procedures. The point of this is partly analytical: by formalizing and studying simple artificial legal procedures, insight can be obtained as to the properties of actual legal procedures. But this also has a normative point. Desirable properties for legal procedures can be precisely formulated, i.e., properties that make procedures fair and/or effective, and whether these procedures have these properties can be formally or computationally investigated. For example, a fairness property is that if the participants' joint knowledge bases logically (according to an argumentation logic) imply a certain outcome, then the procedure allows a dialogue in which this outcome is reached. Insights thus obtained about artificial procedures can then be used in (re-)designing actual legal procedures.

An influential computational model of a legal procedure was Gordon (1994)'s Pleadings Game, which formalized a normative model of pleading, founded on first principles. It was meant to identify the issues to be decided at trial, given what the parties had claimed, conceded, challenged and denied in the pleadings phase and what (defeasibly) follows from it. Other games (usually two-party) define the outcome in terms of whether the adversaries in the end agree on the main issue, while Prakken (2008) defines a three-party game where in the end an adjudicator (for example, a judge or a jury) decides the dispute.

The dynamic and multi-party setting of legal procedures also raises issues of strategy and choice. One of the few investigations of these issues in a legal context is by Riveret et al. (2008), who apply a combination of game theory and argumentation logic to the problem of determining optimal strategies in debates with an adjudicator. In such debates the opposing parties must estimate how likely it is that the premises of their arguments will be accepted by the adjudicator. Moreover, they may have preferences over the outcome of a debate, so that optimal strategies are determined by two factors: the probability of acceptance of their arguments' premises by the adjudicator and the costs/benefits of such arguments.

Conclusion

Computational models of legal reasoning aim to respect that the central notion in legal reasoning is not deduction but argument. Deduction has its place in legal reasoning, but only as part of a larger model of constructing, attacking and comparing arguments. Moreover, computational models of legal reasoning aim to respect that rules are not the only source of legal knowledge: the roles of cases, principles, purposes and values should not be ignored, and these models stress the importance of dynamics, procedure and multi-party interaction. While most work has addressed legal interpretation and normative determinations, some work is addressing legal proof. Here too, deduction is just one of the tools, as part of either probabilistic, argumentation-based or narrative approaches. A main concern here is developing models of legal proof that are rationally well-founded but respect the practical and cognitive constraints faced by the participants in court proceedings.

References

Aleven, V. (2003). Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment, *Artificial Intelligence* **150**: 183–237.

Ashley, K. (1989). Toward a computational theory of arguing with precedents: accomodating multiple interpretations of cases, *Proceedings of the Second International Conference on Artificial Intelligence and Law*, ACM Press, New York, pp. 39–102.

Ashley, K. (1990). *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*, MIT Press, Cambridge, MA.

Atkinson, K., Bench-Capon, T. and McBurney, P. (2005). Arguing about cases as practical reasoning, *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*, ACM Press, New York, pp. 35–44.

Bench-Capon, T. (2002). The missing link revisited: the role of teleology in representing legal argument, *Artificial Intelligence and Law* **10**: 79–94.

Bench-Capon, T. and Coenen, F. (1992). Isomorphism and legal knowledge based systems, *Artificial Intelligence and Law* **1**: 65–86.

Bench-Capon, T. and Sartor, G. (2003). A model of legal reasoning with cases incorporating theories and values, *Artificial Intelligence* **150**: 97–143.

Bennett, W. and Feldman, M. (1981). *Reconstructing Reality in the Courtroom: Justice and Judgment in American Culture*, Methuen-Tavistock, London.

Berman, D. and Hafner, C. (1993). Representing teleological structure in case-based legal reasoning: the missing link, *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, ACM Press, New York, pp. 50–59.

Bex, F. (2011). *Arguments, Stories and Criminal Evidence: A Formal Hybrid Theory*, Springer, Dordrecht.

Bex, F., Prakken, H., Reed, C. and Walton, D. (2003). Towards a formal account of reasoning about evidence: argumentation schemes and generalisations, *Artificial Intelligence and Law* **12**: 125–165.

Branting, L. (1999). *Reasoning with Rules and Precedents: A Computational Model of Legal Analysis*, Kluwer Academic Publishers, Dordrecht/Boston/London.

Chorley, A. and Bench-Capon, T. (2005). An empirical investigation of reasoning with legal cases through theory construction and application, *Artificial Intelligence and Law* **13**: 323–371.

Fenton, N. and Neil, M. (2011). Avoiding legal fallacies in practice using Bayesian networks, *Australian Journal of Legal Philosophy* **36**: 114–151.

Gardner, A. (1987). *An Artificial Intelligence Approach to Legal Reasoning*, MIT Press, Cambridge, MA.

Gordon, T. (1991). An abductive theory of legal issues., *International Journal of Man-Machine Studies* **35**: 95–118.

Gordon, T. (1994). The Pleadings Game: an exercise in computational dialectics, *Artificial Intelligence and Law* **2**: 239–292.

Gordon, T. and Walton, D. (2009). Proof burdens and standards, *in* I. Rahwan and G. Simari (eds), *Argumentation in Artificial Intelligence*, Springer, Berlin, pp. 239–258.

Horty, J. (2011). Rules and reasons in the theory of precedent, *Legal Theory* **17**: 1–33.

Loui, R. and Norman, J. (1995). Rationales and argument moves, *Artificial Intelligence and Law* **3**: 159–189.

McCarty, L. (1977). Reflections on TAXMAN: An experiment in artificial intelligence and legal reasoning, *Harvard Law Review* **90**: 89–116.

McCarty, L. (1995). An implementation of Eisner v. Macomber, *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, ACM Press, New York, pp. 276–286.

Prakken, H. (2008). A formal model of adjudication dialogues, *Artificial Intelligence and Law* **16**: 305–328.

Prakken, H. and Sartor, G. (1996). A dialectical model of assessing conflicting arguments in legal reasoning, *Artificial Intelligence and Law* **4**: 331–368.

Prakken, H. and Sartor, G. (1998). Modelling reasoning with precedents in a formal dialogue game, *Artificial Intelligence and Law* **6**: 231–287.

Prakken, H. and Sartor, G. (2002). The role of logic in computational models of legal argument: a critical survey, *in* A. Kakas and F. Sadri (eds), *Logic Programming and Beyond. Essays In Honour of Robert A. Kowalski, Part II*, number 2048 in *Springer Lecture Notes in Computer Science*, Springer Verlag, Berlin, pp. 342–380.

Prakken, H. and Sartor, G. (2009). A logical analysis of burdens of proof, *in* H. Kaptein, H. Prakken and B. Verheij (eds), *Legal Evidence and Proof: Statistics, Stories, Logic*, Ashgate Publishing, Farnham, pp. 223–253.

Prakken, H. and Sartor, G. (2011). On modelling burdens and standards of proof in structured argumentation, *in* K. Atkinson (ed.), *Legal Knowledge and Information Systems. JURIX 2011: The Twenty-fourth Annual Conference*, IOS Press, Amsterdam etc., pp. 83–92.

Prakken, H. and Sartor, G. (2014). Law and logic: a review. Submitted.

Riveret, R., Prakken, H., Rotolo, A. and Sartor, G. (2008). Heuristics in argumentation: a game-theoretical investigation, *in* P. Besnard, S. Doutre and A. Hunter (eds), *Computational Models of Argument. Proceedings of COMMA 2008*, IOS Press, Amsterdam etc, pp. 324–335.

Rotolo, A. and Governatori, G. (2009). Changing legal systems: Legal abrogations and annulments in defeasible logic, *Logic Journal of IGPL* **18**: 157–94.

Sergot, M. (2001). A computational theory of normative positions, *ACM Transactions on Computational Logic* **2**: 581–622.

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

Skalak, D. and Rissland, E. (1992). Arguments and cases. an inevitable intertwining, *Artificial Intelligence and Law* **1**: 3–44.

Svensson, J. (2002). The use of legal expert systems in administrative decision making, *in* A. Grönlund (ed.), *Electronic Government: Design, Applications and Management*, Idea Group Publishing, London etc, pp. 151–169.

Thagard, P. (2004). Causal inference in legal decision making: Explanatory coherence vs. Bayesian networks, *Applied Artificial Intelligence* **18**: 231–249.

Verheij, B., Hage, J. and van der Herik, H. (1998). An integrated view on rules and principles, *Artificial Intelligence and Law* **6**: 3–26.

Wigmore, J. (1931). *The Principles of Judicial Proof*, 2nd edn, Little, Brown and Company, Boston.