Chapter 1
Argumentation in legal reasoning

Trevor Bench-Capon and Henry Prakken and Giovanni Sartor

1.1 Introduction

A popular view of what Artificial Intelligence can do for lawyers is that it can do no more than deduce the consequences from a precisely stated set of facts and legal rules. This immediately makes many lawyers sceptical about the usefulness of such systems: this mechanical approach seems to leave out most of what is important in legal reasoning. A case does not appear as a set of facts, but rather as a story told by a client. For example, a man may come to his lawyer saying that he had developed an innovative product while working for Company A. Now Company B has made him an offer of a job, to develop a similar product for them. Can he do this? The lawyer firstly must interpret this story, in the context, so that it can be made to fit the framework of applicable law. Several interpretations may be possible. In our example it could be seen as being governed by his contract of employment, or as an issue in Trade Secrets law. Next the legal issues must be identified and the pros and cons of the various interpretations considered with respect to them. Does his contract include a non-disclosure agreement? If so, what are its terms? Was he the sole developer of the product? Did Company A support its development? Does the product use commonly known techniques? Did Company A take measures to protect the secret? Some of these will favour the client, some the Company. Each interpretation will require further facts to be obtained. For example, do the facts support a claim that the employee was the sole developer of the product? Was development work carried

Trevor Bench-Capon
Department of Computer Science, University of Liverpool, UK e-mail: tbc@liverpool.ac.uk

Henry Prakken
Department of Information and Computing Sciences, Utrecht University, and Faculty of Law, University of Groningen, The Netherlands e-mail: henry@cs.uu.nl

Giovanni Sartor
European University Institute, Law Department, Florence, and CIRSFID, University of Bologna, Italy e-mail: giovanni.sartor@eui.eu

out in his spare time? What is the precise nature of the agreements entered into? Once an interpretation has been selected, the argument must be organised into the form considered most likely to persuade, both to advocate the client's position and to rebut anticipated objections. Some precedents may point to one result and others to another. In that case, further arguments may be produced to suggest following the favourable precedent and ignoring the unfavourable one. Or the rhetorical presentation of the facts may prompt one interpretation rather than the other. Surely all this requires the skill, experience and judgement of a human being? Granted that this is true, much effort has been made to design computer programs that will help people in these tasks, and it is the purpose of this chapter to describe the progress that has been made in modelling and supporting this kind of sophisticated legal reasoning.

We will review[1] systems that can store conflicting interpretations and that can propose alternative solutions to a case based on these interpretations. We will also describe systems that can use legal precedents to generate arguments by drawing analogies to or distinguishing precedents. We will discuss systems that can argue why a rule should not be applied to a case even though all its conditions are met. Then there are systems that can act as a mediator between disputing parties by structuring and recording their arguments and responses. Finally we look at systems that suggest mechanisms and tactics for forming arguments.

Much of the work described here is still research: the implemented systems are prototypes rather than finished systems, and much work has not yet reached the stage of a computer program but is stated as a formal theory. Our aim is therefore to give a flavour (certainly not a complete survey) of the variety of research that is going on and the applications that might result in the not too distant future. Also for this reason we will informally paraphrase example inputs and outputs of systems rather than displaying them in their actual, machine-readable format; moreover, because of space limitations the examples have to be kept simple.

1.2 Early systems for legal argumentation

In this section we briefly discuss some early landmark systems for legal argumentation. All of them concern the construction of arguments and counterarguments.

1.2.1 Conflicting interpretations

Systems to address conflicting interpretations of legal concepts go back to the very beginnings of AI and Law. Thorne McCarty (e.g. (25; 27)) took as his key problem a landmark Supreme Court Case in US tax law which turned on differing interpretations of the concept of ownership, and set himself the ambitious goal of reproducing,

---

[1] This chapter is a revised and updated version of (6).

in his TAXMAN system, both the majority and the dissenting opinions expressing these interpretations. This required highly sophisticated reasoning, constructing competing theories and reasoning about the deep structure of legal concepts to map the specific situation onto paradigmatic cases. Although some aspects of the system were prototyped, the aim was perhaps too ambitious to result in a working system, certainly given the then current state of the art. This was not McCarty's goal, however: his motivation was to gain insight into legal reasoning through a computational model. McCarty's main contribution was the recognition that legal argument involves theory construction as well as reasoning with established knowledge. McCarty (26) summarises his position as follows: "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" (p. 285). Note also the emphasis on persuasion, indicating that we should expect to see argumentation rather than proof. Both the importance of theory construction and the centrality of persuasive argument are still very much part of current thinking in AI and Law.

Another early system was developed by Anne Gardner (15) in the field of offer and acceptance in American contract law. The task of the system was "to spot issues": given an input case, it had to determine which legal questions arising in the case were easy and which were hard, and to solve the easy ones. The system was essentially rule based, and this simpler approach offered more possibilities for practical exploitation than did McCarty's system. One set of rules was derived from the Restatement of Contract Law, a set of 385 principles abstracting from thousands of contract cases. These rules were intended to be coherent, and to yield a single answer if applicable. This set of rules was supplemented by a set of interpretation rules derived from case law, common sense and expert opinion, intended to link these other rules to the facts of the case. Gardner's main idea was that easy questions were those where a single answer resulted from applying these two rule sets, and hard questions, or issues, were either those where no answer could be produced, because no interpretation rule linked the facts to the substantive rules, or where conflicting answers were produced by the facts matching with several rules. Some of the issues were resolved by the program with a heuristic that gives priority to rules derived from case law over restatement and commonsense rules. The rationale of this heuristic is that if a precedent conflicts with a rule from another source, this is usually because that rule was set aside for some reason by the court. The remaining issues were left to the user for resolution.

Consider the following example, which is a very much simplified and adapted version of Gardner's own main example[2]. The main restatement rule is

R1: An offer and an acceptance constitute a contract

Suppose further that there are the following commonsense (C) and expert (E) rules on the interpretation of the concepts of offer and acceptance:

---

[2] We in particular abstract from Gardner's refined method for representing knowledge about (speech act) events.

C1: A statement "Will supply ..." in reply to a request for offer is an offer.
C2: A statement "Will you supply ..." is a request for offer.
C3: A statement "I accept ..." is an acceptance.
E1: A statement "I accept" followed by terms that do not match the terms of the offer is not an acceptance.

Suppose that Buyer sent a telegram to Seller with "Will you supply carload salt at $2.40 per cwt?" to which Seller replied with "Will supply carload at $2.40, terms cash on delivery", after which Buyer replied with her standard "Purchase Order" indicating "I accept your offer of 12 July" but which also contained a standard provision "payment not due until 30 days following delivery".

Applying the rules to these events, the "offer" antecedent of R1 can be established by C1 combined with C2, since there are no conflicting rules on this issue. However, with respect to the "acceptance" antecedent of R1 two conflicting rules apply, viz. C3 and E1. Since we have no way of giving precedence to C3 or E1, the case will be a hard one, as there are two conflicting notions of "acceptance". If the case is tried and E1 is held to have precedence, E1 will now be a precedent rule, and any subsequent case in which this conflict arises will be easy, since, as a precedent rule, E1 will have priority over C3.

1.2.2 Reasoning with precedents

The systems described in the last section do recognise the importance of precedent cases as a source of legal knowledge, but they make use of them by extracting the rationale of the case and encoding it as a rule. To be applicable to a new case, however, the rule extracted may need to be analogised or transformed to match the new facts. Nor is extracting the rationale straightforward: judges often leave their reasoning implicit and in reconstructing the rationale a judge could have had in mind there may be several candidate rationales, and they can be expressed at a variety of levels of abstraction. These problems occur especially in so-called 'factor-based' domains, i.e., domains where problems are solved by considering a variety of factors that plead for or against a solution. In such domains a rationale of a case often just expresses the resolution of a particular set of factors in a specific case. A main source of conflict in such domains is that a new case often will not exactly match a precedent but will share some features with it, lack some of its other features, and/or have some additional features. Moreover, cases are more than simple rationales: matters such as the context and the procedural setting can influence the way the case should be used. In consequence, some researchers have attempted to avoid using rules and rationales altogether, instead representing the input, often interpreted as a set of factors, and the decisions of cases, and defining separate argument moves for interpreting the relation between the input and decision (e.g. (23; 1), both to be discussed below). This approach is particularly associated with researchers in US, where the common law tradition places a greater stress on precedent cases and their particular features than is the case with the civil law jurisdictions of Europe. None

the less cases are also used in civil law jurisdictions and the reasoning techniques are similar.

The most influential system of this sort is HYPO (2), developed by Edwina Rissland and Kevin Ashley in the domain of US Trade Secrets Law, which can be seen as a factor-based domain. In HYPO cases are represented according to a number of dimensions. A dimension is some aspect of the case relevant to the decision, for example, the security measures taken by the plaintiff. One end of the dimension represents the most favourable position for the plaintiff (e.g. specific non-disclosure agreements), while the other end represents the position most favourable to the defendant (e.g. no security measures at all). Typically a case will lie somewhere between the two extremes and will be more or less favourable accordingly. HYPO then uses these dimensions to construct three-ply arguments. First one party (say the plaintiff) cites a precedent case decided for that side and offers the dimensions it shares with the current case as a reason to decide the current case for that side. In the second ply the other party responds either by citing a counter example, a case decided for the other side which shares a different set of dimensions with the current case, or distinguishing the precedent by pointing to features which make the precedent more, or the current case less, favourable to the original side. In the third ply the original party attempts to rebut the arguments of the second ply, by distinguishing the counter examples, or by citing additional precedents to emphasise the strengths or discount the weaknesses in the original argument.

Subsequently Ashley went on, with Vincent Aleven, to develop CATO (most fully reported in (1)), a system designed to help law students to learn to reason with precedents. CATO simplifies HYPO in some respects but extends it in others. In CATO the notion of dimensions is simplified to a notion of factors. A factor can be seen as a specific point of the dimension: it is simply present or absent from a case, rather than present to some degree, and it always favours either the plaintiff or defendant. A new feature of CATO is that these factors are organised into a hierarchy of increasingly abstract factors, so that several different factors can be seen as meaning that the same abstract factor is present. One such abstract factor is that the defendant used questionable means to obtain the information, and two more specific factors indicating the presence of this factor are that the defendant deceived the plaintiff and that he bribed an employee of the plaintiff: both these factors favour the plaintiff. The hierarchy allows for argument moves that interpret the relation between a case's input and its decision, such as emphasising or downplaying distinctions. To give an example of downplaying, if in the precedent defendant used deception while in the new case instead defendant bribed an employee, then a distinction made by the defendant at this point can be downplayed by saying that in both cases the defendant used questionable means to obtain the information. To give an example of emphasising a distinction, if in the new case defendant bribed an employee of plaintiff while in the precedent no factor indicating questionable means was present, then the plaintiff can emphasise the distinction "unlike the precedent, defendant bribed an employee of plaintiff" by adding "and therefore, unlike the precedent defendant used questionable means to obtain the information".

Perhaps the most elaborate representation of cases was produced in Karl Branting's Grebe system in the domain of industrial injury, where cases were represented as semantic networks (13). The program matched portions of the network for the new case with parts of the networks of precedents, to identify appropriate analogies.

Of all this work, HYPO in particular was highly influential, both in the explicit stress it put on reasoning with cases as constructing arguments, and in providing a dialectical structure in which these arguments could be expressed, anticipating much other work on dialectical procedures.

1.3 Logical accounts of reasoning under disagreement

The systems discussed in the previous section were (proposals for) implemented systems, based on informal accounts of some underlying theory of reasoning. Other AI & Law research aims at specifying theories of reasoning in a formal way, in order to make general reasoning techniques from nonmonotonic logic and formal argumentation available for implementations.

The first AI & Law proposals in this vein, for example, (16; 30), can be regarded as formal counterparts of Gardner's ideas on issue spotting. Recall that Gardner allows for the presence in the knowledge base of conflicting rules governing the interpretation of legal concepts and that she defines an issue as a problem to which either no rules apply at all, or conflicting rules apply. Now in logical terms an issue can be defined as a proposition such that either there is no argument about this proposition or there are both arguments for the proposition and for its negation.

Some more recent work in this research strand has utilised a very abstract AI framework for representing systems of arguments and their relations developed by Dung (14). For Dung, the notion of argument is entirely abstract: all that can be said of an argument is which other arguments it attacks, and which it is attacked by. Given a set of arguments and the attack relations between them, it is possible to determine which arguments are acceptable: an argument which is not attacked will be acceptable, but if an argument has attackers it is acceptable only if it can be defended, against these attackers, by acceptable arguments which in turn attack those attackers. This framework has proved a fruitful tool for understanding nonmonotonic logics and their computational properties. Dung's framework has also been made use of in AI & Law. It was first applied to the legal domain by Prakken & Sartor (35), who defined a logic for reasoning with conflicting rules as an instantiation of Dung's framework. (See below and Chapter **??** of this book) Bench-Capon (3) has explored the potential of the fully abstract version of the framework to represent a body of case law. he uses preferred semantics, where arguments can defend themselves: in case of mutual attack this gives rise to multiple sets of acceptable arguments, which can explain differences in the application of law in different jurisdictions, or at different times in terms of social choices. Dung's framework has also been extended to include a more formal consideration of social values (see Chapter **??** of this book).

1.3.1 Reasoning about conflicting rules

Generally speaking, the proposed systems discussed so far attempt to identify conflicting interpretations and arguments, but do not attempt to resolve them, leaving it to the user to choose which argument will be accepted. As we saw above, Gardner's system went somewhat further in that it gave priority to rules derived from case law over restatement and commonsense rules. Thus her system was able to solve some of the cases to which conflicting rules apply. This relates to much logical work in Artificial Intelligence devoted to the resolution of rule conflicts in so-called commonsense reasoning. If we have a rule that birds can fly and another that ostriches cannot fly, we do not want to let the user decide whether Cyril the ostrich can fly or not: we want the system to say that he cannot, since an ostrich is a specific kind of bird. Naturally attempts have been made to apply these ideas to law.

One approach was to identify general principles used in legal systems to establish which of two conflicting rules should be given priority. These principles included preferring the more specific rule (as in the case of the ostrich above, or where a law expresses an exception to a general provision), preferring the more recent rule, or preferring the rule deriving from the higher legislative authority (for instance, 'federal law precedes state law'). To this end the logics discussed above were extended with the means to express priority relations between rules in terms of these principles so that rule conflicts would be resolved. Researchers soon realised, however, that general priority principles can only solve a minority of cases. Firstly, as for the specificity principle, whether one rule is more specific than another often depends on substantive legal issues such as the goals of the legislator, so that the specificity principle cannot be applied without an intelligent appreciation of the particular issue. Secondly, general priority principles usually only apply to rules from regulations and not to, for instance, case rationales or interpretation rules derived from cases. Accordingly, in many cases the priority of one rule over another can be a matter of debate, especially when the rules that conflict are unwritten rules put forward in the context of a case. For these reasons models of legal argument should allow for arguments about which rule is to be preferred.

As an example of arguments about conflicting case rationales, consider three cases discussed by, amongst others, (10; 7; 32) and (8) concerning the hunting of wild animals. In all three cases, the plaintiff (P) was chasing wild animals, and the defendant (D) interrupted the chase, preventing P from capturing those animals. The issue to be decided is whether or not P has a legal remedy (a right to be compensated for the loss of the game) against D. In the first case, *Pierson v Post*, P was hunting a fox on open land in the traditional manner using horse and hound, when D killed and carried off the fox. In this case P was held to have no right to the fox because he had gained no possession of it. In the second case, *Keeble v Hickeringill*, P owned a pond and made his living by luring wild ducks there with decoys, shooting them, and selling them for food. Out of malice, D used guns to scare the ducks away from the pond. Here P won. In the third case, *Young v Hitchens*, both parties were commercial fisherman. While P was closing his nets, D sped into the gap, spread his

own net and caught the fish. In this case D won. The rules we are concerned with here are the rationales of these cases:

R1 *Pierson*: If the animal has not been caught, the defendant wins
R2 *Keeble*: If the plaintiff is pursuing his livelihood, the plaintiff wins
R3 *Young*:   If the defendant is in competition with the plaintiff and the animal is not caught, the defendant wins.

Note that R1 applies in all cases and R2 in both *Keeble* and *Young*. In order to explain the outcomes of the cases we need to be able to argue that R3 > R2 > R1. To start with, note that if, as in HYPO, we only look at the factual similarities and differences, none of the three precedents can be used to explain the outcome of one of the other precedents. For instance, if we regard *Young* as the current case, then both *Pierson* and *Keeble* can be distinguished. A way of arguing for the desired priorities, first mooted in (10), is to refer to the purpose of the rules, in terms of the social values promoted by following the rules.

The logic of (35) provides the means to formalise such arguments. Consider another case in which only plaintiff was pursuing his livelihood and in which the animal was not caught. In the following (imaginary) dispute the parties reinterpret the precedents in terms of the values promoted by their outcomes, in order to find a controlling precedent (we leave several details implicit for reasons of brevity; a detailed formalisation method can be found in (32); see also (8)).

Plaintiff: I was pursuing my livelihood, so (by *Keeble*) I win
Defendant: You had not yet caught the animal, so (by *Pierson*) I win
Plaintiff: following *Keeble* promotes economic activity, which is why *Keeble* takes precedence over *Pierson*, so I win.
Defendant: following *Pierson* protects legal certainty, which is why *Keeble* does not take precedence over *Pierson*, so you do not win.
Plaintiff: but promoting economic activity is more important than protecting legal certainty since economic development, not legal certainty is the basis of this country's prosperity. Therefore, I am right that *Keeble* takes precedence over *Pierson*, so I still win.

This dispute contains priority debates at two levels: first the parties argue about which case rationale should take precedence (by referring to values advanced by following the rationale), and then they argue about which of the conflicting preference rules for the rationales takes precedence (by referring to the relative order of the values). In general, a priority debate could be taken to any level and will be highly dependent on the context and jurisdiction. Various logics proposed in the AI & Law literature are able to formalise such priority debates, such as (17; 35; 19) and (21).

1.3.2 Other arguments about rules

Besides priority debates in case of conflicting rules, these logics can also model debates about certain properties of rules, such as their legal validity or their applicability to a legal case. The most fully developed logical theory about what it takes to apply a rule is reason-based logic, developed jointly by Jaap Hage and Bart Verheij (e.g. (19; 45)). They claim that applying a legal rule involves much more than subsuming a case under the rule's conditions. Their account of rule application can be briefly summarised as follows. First in three preliminary steps it must be determined whether the rule's conditions are satisfied, whether the rule is legally valid, and whether the rule's applicability is not excluded in the given case by, for instance, a statutory exception. If these questions are answered positively (and all three are open to debate), it must finally be determined that the rule can be applied, i.e., that no conflicting rules or principles apply. On all four questions reason-based logic allows reasons for and against to be provided and then weighed against each other to obtain an answer.

Consider by way of illustration a recent Dutch case (HR 7-12-1990, *NJ* 1991, 593) in which a male nurse aged 37 married a wealthy woman aged 97 whom he had been nursing for several years, and killed her five weeks after the marriage. When the woman's matrimonial estate was divided, the issue arose whether the nurse could retain his share. According to the relevant statutes on Dutch matrimonial law the nurse was entitled to his share since he had been the woman's husband. However, the court refused to apply these statutes, on the grounds that applying it would be manifestly unjust. Let us assume that this was in turn based on the legal principle that no one shall profit form his own wrongdoing (the court did not explicitly state this). In reason-based logic this case could be formalised as follows (again the full details are suppressed for reasons of brevity).

Claimant: Statutory rule R is a valid rule of Dutch law since it was enacted according to the Dutch constitution and never repealed. All its conditions are satisfied in my case, and so it should be applied to my case. The rule entitles me to my late wife's share in the matrimonial estate. Therefore, I am entitled to my wife's share in the matrimonial estate.
Defendant: Applying rule R would allow you to profit from your own wrongdoing: therefore rule R should not be applied in this case.
Court: The reason against applying this rule is stronger than that for applying the rule, and so the rule does not apply.

Of course, in the great majority of cases the validity or applicability of a statute rule is not at issue but instead silently presumed by the parties (recall the difference between arguments and proofs described in the introduction). The new logical techniques alluded to above can also deal with such presumptions, and they can be incorporated in reason-based logic.

Reason-based logic also has a mechanism for 'accruing' different reasons for conclusions into sets and for weighing these sets against similar sets for conflict-

ing conclusions. Thus it captures that having more reasons for a conclusions may strengthen one's position. Prakken (33) formalises a similar mechanism for accruing arguments in the context of Dung's framework. Prakken also proposes three principles that any model of accrual should satisfy. In Chapter **??** of this book Gordon and Walton show that the Carneades logic of (18) has an accrual mechanism that satisfies these principles.

One way to argue about the priority of arguments is to claim that the argument is preferred if it is grounded in the better or more coherent legal theory[3]. While there has been considerable progress in seeing how theories can be constructed on the basis of a body of past cases, evaluation of the resulting theories in terms of their coherence is more problematic, since coherence is a difficult notion to define precisely[4]. Bench-Capon & Sartor (8) describe some features of a theory which could be used in evaluation, such as simplicity of a theory or the number of precedent cases explained by the theory. As an (admittedly somewhat simplistic) example of the last criterion, consider again the three cases on hunting animals, and imagine two theories that explain the case decisions in terms of the values of promotion of economic activity and protection of legal certainty. A theory that gives precedence to promoting economic activity over protecting legal certainty explains all three precedents while a theory with the reverse value preference fails to explain *Keeble*. The first theory is therefore on this criterion the more coherent one. However, how several coherence criteria are to be combined is a matter for further research. For an attempt to give a metric for coherence, see (7). Coherence is also discussed in (20), where coherence is treated mainly in terms of respecting *a fortiori* arguments.

1.4 Dialogue and mediation systems

Implicit in the notion of argument is that there are two parties with opposing views. Already in HYPO there is the dialectical structure of point, counter point and rebuttal, and most logics for argumentation discussed above also have this dialectical flavour. It is therefore a natural step to make this dialogical structure explicit, and to build systems to conduct or mediate dialogues between the opposed parties. Such dialogue systems also provide the opportunity to model the procedure under which a dispute is conducted, and the context in which information is introduced to a dispute. Taking a procedural point of view forces us to think about matters such as burden of proof, admissibility of evidence, agreed and contested points, and the role of a neutral third party to arbitrate the dispute.

One of the first such systems in AI & Law was Tom Gordon's Pleadings Game, which embodies an idealised model of civil pleadings in common law systems (17). The objective of the system is to extend the issue-spotting task of Gardner's pro-

---

[3] There is, of course, a debate in legal theory as to how we can provide an epistemology of law, and coherence is only one position. Coherence is discussed here as it is the position which has received most attention in AI & Law

[4] For fuller discussions of coherence, see (29) and (28, Ch. 2)

gram to a dialogical setting. It is to allow two human parties to state the arguments and facts that they believe to be relevant, so that they can determine where they agree and where they disagree. The residual disagreements will go on to form the issues when the case is tried. The system plays two roles in this process: it acts as a referee to ensure that the proper procedure is followed, and records the facts and arguments that are presented and what points are disputed, so as to identify the issues that require resolution. The Pleadings Game has a built-in proof mechanism for an argumentation logic, which is applied to check the logical well-formedness of the arguments stated by the user, and to compute which of the stated arguments prevail, on the basis of the priority arguments also stated by the user and a built-in specificity checker. The main addition to Gardner's system is that in the Pleadings Game not only the content of the arguments is relevant but also the attitudes of the parties expressed towards the arguments and their premises.

Let us illustrate this with the following simplified dispute, based on the example that we above used to illustrate Gardner's system.

Plaintiff: I claim (1) we have a contract
Defendant: I deny 1
Plaintiff: We have a valid contract since (2) I made an offer and (3) you accepted it, so we have a contract.
Defendant: I concede 2 but I deny 3.
Plaintiff: (4) you said "I accept...", so by C1 you accepted my offer.
Defendant: I concede 4 and C1, but (5) my statement "I accept ..." was followed by terms that do not match the terms of your offer. So by P1 (which takes priority over C1) I (6) did not accept you offer.
Plaintiff: I concede P1 and that P1 takes priority over C1 but I deny 5.
Defendant: (7) you required payment upon delivery while (8) I offered payment 30 days following delivery, so there is a mismatch between our terms.
Plaintiff: I concede (7) and the argument but I deny (8).

At this point, there is one argument for the conclusion that a contract was created, based on the premises 2, 4 and C1 (note that plaintiff left R1 implicit and defendant silently agreed with this). The intermediate conclusion (3) of this argument that there was an acceptance is defeated by a counterargument based on premises 7, 8 and P1. So according to a purely logical analysis of the dispute the case is easy, having as outcome that no contract exists between the parties. This agrees with Gardner's treatment of the example. However, in the Pleadings Game it also matters that the plaintiff has denied defendant's claim (8). This is a factual issue making the case hard, and which has to be decided in court.

The Pleadings Game was fully implemented, but purely as an experimental system: in particular the arguments had to be presented in a complicated logical syntax so that they could be handled by the underlying proof mechanism. The trade-off between ease of use and the ability of the system to process the information it receives remains a difficult problem for such systems.

Following Gordon's work, a number of other systems for dialogue were produced. Arno Lodder's Dialaw (22) is a two-player dialogue game that combines the notion of propositional commitment (48) with Hage and Verheij's Reason Based Logic. The players can use locutions for claiming a proposition and for challenging, conceding and retracting a claimed proposition. Arguments are constructed implicitly, by making a new claim in reply to a challenge. Arguments can also be about the procedural correctness of dialogue moves. Each dialogue begins with a claim of one player, and then the turn usually switches after each move. When the commitments of one player logically imply a claim of the other player, the first player must either concede it or retract one of the implying commitments. A dialogue terminates if no disagreement remains, i.e., if no commitment of one player is not also a commitment of the other. The first player wins if at termination he is still committed to his initial claim, the second player wins otherwise.

Bench-Capon's TDG (5) is intended to produce more natural dialogues than the "stilted" ones produced by systems such as the Pleadings Game and Dialaw. To this end, its speech acts are based on Toulmin's well-known argument scheme (44). In this scheme, a *claim* is supported by *data*, which support is *warranted* by an inference licence, which is *backed* by grounds for its acceptance; finally, a claim can be attacked with a *rebuttal*, which itself is a claim and thus the starting point of a counterargument. Arguments can be chained by regarding data also as claims, for which data can in turn be provided. TDG has speech acts for asking for and providing these elements of an argument; a dialogue starts with a claim and then the protocol supports a dialogue which constructs a Toulmin structure whilst subjecting it to a top-down critical examination.

In (34) Prakken proposes a dialogue game model for "adjudication dialogues", by adding a neutral third party who adjudicates the conflict between two adversaries. The main feature of the model is a division into an argumentation phase, where the adversaries plea their case and the adjudicator has a largely mediating role, and a decision phase, where the adjudicator decides the dispute on the basis of the claims, arguments and evidence put forward in the argumentation phase. The model allows for explicit decisions on admissibility of evidence and burden of proof by the adjudicator in the argumentation phase. Adjudication is modelled as putting forward arguments, in particular undercutting and priority arguments, in the decision phase.

Such a dialogue game model of adjudication paves the way for formal models of burden of proof and presumption. Research on legal argumentation has recently addressed these issues, which has required models going beyond existing 'standard' non-monotonic logics. In (31) the logic of (35) is adapted to make the acceptability of an argument dependent on how such a burden is distributed over the parties, while in the original system it is always allocated to one side. In (36) the resulting logic is extended with the means to explicitly represent and reason about distributions of the burden of proof. They also claim that if such a logic is embedded in a dialogue model for adjudication, a precise distinction can be made between three kinds of burdens of proof often distinguished in the law, namely, the burden of production, the burden of persuasion and the tactical burden of proof. The resulting combination of logic and dialogue game is applied to presumptions in (37): by developing the idea that legal

presumptions can be modelled as defeasible rule, they provide ways of representing how a legal presumption can be defeated by evidence to the contrary and how the burden of proof can be allocated with regard to the disproval of presumed facts. In related work, Gordon, Prakken & Walton (18) have provided a logical model of reasoning under standards of proof within the Carneades logic. In Chapter **??** of this volume, Gordon and Walton extend it with definitions of these three and some further notions of burden of proof and compare their account with the work of Prakken & Sartor.

1.5 Tactics for dispute

Once arguments are placed in a dialogical setting, it becomes apparent that at various points of the dialogue, the parties will have a choice of moves by which to attack their opponent or defend their own arguments. Questions then arise as to which moves are available to construct, attack and defend arguments, and whether there are principles to guide the choice of move. In fact, the implemented dialogue systems of the previous section do not address these questions, because they are intended to act as a mediator between two human players. The responsibility of the system is thus limited to enforcing the rules of the game, while strategy and tactics are the responsibility of the human users.

In their work on the CABARET system, Skalak & Rissland (43) attempted to identify arguments that could be made in a dispute using rules and cases. They begin by identifying a number of forms of argument, and then describe argument strategies to be used according to the context of the dispute. For example, if the current case matches with most but not all the features of some statutory rule that one wishes to use, the rule must be broadened so as to make the rule applicable to the case. Or if a rule is applicable to the case under consideration but would be unfavourable, that rule needs to be discredited. They then identify the moves that can be made to realise the strategies, depending on the disposition of the precedent, and whether the precedent does or does not establish the desired consequent. One move to broaden a rule is to find a precedent that also lacked the missing features but in which the conclusion of the rule was nevertheless drawn. To discredit a rule one can try to find a precedent case in which it was not followed even though all its conditions were satisfied in the case. Finally they identify a number of primitive operations in terms of which the moves can be realised. These operations include all moves that can be made in HYPO with cases. All of this is then brought together in a decision tree which suggests which strategy should be adopted, which moves need to be used to fulfill it and which primitives will enable the required moves.

Loui & Norman (23) take this approach a step further in their formal model of the use of rationales in disputes. They allow for a position under attack to be first restated, in order to make the attack more effective. For example if an argument using a rationale *if P then Q* is to be attacked, it may be helpful to restate this as

if *P then R and if R then Q*, and to provide a counter example to *if P then R*. They provide a number of other examples of rationales and tactics for attacking them.

CABARET, by distinguishing different kinds of building materials, and providing different moves and attacks appropriate to each kind, can produce its elegant classification of strategies. The central idea of distinguishing different kinds of premises and different ways of dealing with them is explicitly addressed by work on argument schemes, which we discuss in the next section.

Finally, in (41) and (40) strategic aspects of legal argument are modelled with game-theoretic methods. A game-theoretical analysis of legal argumentation requires a method of determining the payoffs the parties will receive as a consequence of the arguments they present in a dispute (combined with the choices of the other parties) as well as the identification of the equilibrium strategies the parties should pursue (namely what strategies of the two parties are such that each one of them represents a best response to the strategy of the other). The first aspect has been addressed in (41), while the second has been considered in (40).

1.6 Argument schemes

In a logical proof we have a set of premises and a conclusion which is said to follow from them. The premises are considered to be entirely homogeneous. Many of the systems discussed so far likewise make no distinctions among their premises. In arguments expressed in natural language in contrast we can typically see the premises as playing different roles in the argument. By identifying these roles, we can present the arguments in a more readily understandable fashion, and also identify the various different ways in which the argument may be attacked. Structuring the argument in this way produces an argument scheme. Analysing legal reasoning in terms of argument schemes produces a taxonomy of arguments, which may provide useful guidance for building implemented argumentation systems, analogous to the guidance provided by domain ontologies for building knowledge-based systems (cf. e.g. (28)).

One argument scheme that has been widely used in AI and Law is that devised by Toulmin (44) (see the description of TDG in Section 1.4)). This has been mainly used to present arguments to users, as in PLAID (9) and SPLIT UP (52), but it has also been used as the basis of a dialogue game, Bench-Capon's TDG, in which the moves of the game relate to providing various elements of the scheme.

While Toulmin attempts to supply a general scheme for arguments, others have attempted to classify arguments in terms of various specific schemes (47). One of the schemes discussed by Walton (pp. 61–63) is the scheme of arguments from the position to know:

Person *W* says that *p*
Person *W* is in the position to know about *p*
Therefore, *p*

Walton also discusses two special versions of this scheme for witness and expert testimonies. Clearly, these schemes are very relevant for evidential legal reasoning. Another scheme discussed by Walton (pp. 75–77) is the scheme from good (or bad) consequences:

> If *A* is brought about, then good (bad) consequences will (may plausibly) occur. *Therefore*, *A* should (not) be brought about.

One instantiation is adapted from a recent discussion in Dutch privacy law whether email addresses are personal data.

> If the term "personal data" of the Dutch Data Protection Act is interpreted to include email addresses, then legal measures against spam become possible, which is good. *Therefore*, the term "personal data" of the Dutch Data Protection Act should be interpreted to include email addresses.

Argument schemes are not classified according to their logical form but according to their content. Many argument schemes in fact express epistemological principles (such as the scheme from the position to know) or principles of practical reasoning (such as the scheme from consequences). Accordingly, different domains may have different sets of such principles. Each argument scheme comes with a customised set of critical questions that have to be answered when assessing whether their application in a specific case is warranted. Thus with argument schemes it becomes clear that the different premises are each associated with their own particular types of attack, in contrast to the purely logical systems in which attacks are uniform. Some of these questions pertain to acceptability of the premises, such as "is *W* in the position to know about *p*?" or "is the possibility to use legal means against spam really good?". Other critical questions point at exceptional circumstances in which the scheme may not apply, such as "is *W* sincere?" or "are there better ways to bring about these good consequences?". Clearly, the possibility to ask such critical questions makes argument schemes defeasible, since negative answers to such critical questions are in fact counterarguments, such as "Person *W* is not sincere since he is a relative of the suspect and relatives of suspects tend to protect the suspect". Another reason why argument schemes are defeasible is that they may be contradicted by conflicting applications of the same or another scheme. For instance, a positive instance of the scheme from consequences can be attacked by a negative instance of the same scheme, such as by "interpreting email addresses as personal data also has bad consequences, since the legal system will be flooded with litigation, so the term "personal data" should not be interpreted to include email addresses". Or one person in a position to know (say an eyewitness) may have said that the suspect was at the crime scene while another eyewitness may have said that the suspect was not at the crime scene.

Originally only Toulmin's argumentation scheme was paid explicit attention in AI & Law, although implicit appeal to other argumentation schemes can be seen in many of the systems discussed above. For example, HYPO identifies the two ways in which the citation of a precedent may be attacked, and reason-based logic identifies ways to reason about the application of legal rules. More recently explicit argumentation schemes have been used. For example, (4) use an extended version

of the scheme from consequences to link legal arguments and social values, and this approach has been applied to the representation of a particular case by (51). (50) presented a set of particular argumentation schemes designed to capture the cased based reasoning used in CATO. (12) modelled several schemes for reasoning about evidence, and this work has subsequently been developed by (11). In Carneades (18) a generalised notion of argument schemes is applied to legal cases; Carneades is further described in chapter **??** of this volume.

1.7 Systems to structure argument

Arguments can often be rather complex, so that understanding the web of relationships becomes difficult. There is clear potential for computers to provide a means of addressing this problem. The idea of providing a visual means of structuring legal arguments is not new to the legal field: as early as the 1930s John Henry Wigmore (49) produced a graphical notation for depicting legal arguments and their relations of support and attack, so as to make sense of a mass of evidence. In this way the relationships between the evidence and the point to be proven, and the ways in which the chain of reasoning could be attacked could be clearly seen.

In Wigmore's days the only way to draw such graphs was with pencil and paper, which perhaps explains why his method was forgotten until Schum & Tillers (42) saw the potential of the computer for supporting the drawing and manipulation of such graphs. They proposed a software system MarshalPlan for visualising preliminary fact investigation based on Wigmore's diagrams. Some other systems within AI & Law that provide support for the graphical structuring of argumentation are Verheij's ArguMed system (46) and Loui's Room 5 system (24), which replaces ArguMed's (and MarshalPlan's) graph structures with encapsulated text boxes, to avoid "pointer spaghetti". The Araucaria system (38) combines an ArguMed-like graphical notation with means to label the arguments as instances of predefined argumentation schemes, which are stored in a database together with their critical questions. The Avers system (11) in turns combines an Araucaria-like approach with the means to visualise crime scenarios as abductive reasoning. An overview of various argument visualisation tools and their legal applications is given in (39).

Argument structuring systems have uses in areas where the clear presentation of the argument is of prime importance. They could be used in preliminary fact investigation (see MarshalPlan or Avers), in teaching (many argument structuring systems outside the legal domain have been developed especially for teaching), for case management or for mediation in online dispute resolution. In all these cases, the usefulness of such systems might be increased by integrating them with documentary sources. For instance, when supporting preliminary fact investigation, the structured evidential arguments could be linked to police documents containing the available evidence. Or when used for case management, the structured arguments could be linked to the case files. Or when a structuring system is used for teaching the analysis of a case decision, the structured arguments could be linked to

the corresponding fragment in the case decisions in the casebook used by the students. Work on argumentation schemes can further augment the usefulness of such systems. When constructing arguments, argument schemes provide a repertoire of forms of argument to be considered, and a template prompting for the pieces that are needed; when attacking arguments they provide a set of critical questions that can identify potential weaknesses in the opponents case. Araucaria and Avers provide examples of research systems pointing in this direction.

1.8 Concluding remarks

In this chapter we have given an overview of the ways in which argumentation has been addressed in legal applications. Legal reasoning has many distinctive features, which include: any proposed set of rules inevitably contain gaps and conflicts; many of its concepts are imprecisely defined meaning that interpretation is required; precedent cases play an important role; procedural issues can influence the status of arguments; much legal argumentation is adversarial and dialectic in nature; the facts of a case need to be selected and characterised; many decisions express a preference for particular values and purposes; and all its conclusions are defeasible, subject often to formal appeal. All of these features mean that deduction cannot provide an adequate model of legal reasoning and instead argumentation must take centre stage to allow for these contextual, procedural and interpretative elements. For this reason developments in computational models of argumentation have been readily taken up by the AI & Law community. Equally, however, the legal AI community has contributed much to computational models of argumentation: a considerable amount of the work described in this book has its origins in work motivated by legal applications, and more than half the chapters have authors who have published in specialist AI & Law venues.

The legal domain can thus act both a motivation and a test-bed for developments in argumentation, and we look forward to much fruitful future interaction between the two areas.

References

[1] V. Aleven. *Teaching Case-Based Argumentation Through a Model and Examples*. PhD Dissertation University of Pittsburgh, 1997.
[2] K. Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. MIT Press, Cambridge, MA, 1990.
[3] T. Bench-Capon. Representation of case law as an argumentation framework. In *Legal Knowledge and Information Systems. JURIX 2002: The Fifteenth Annual Conference*, pages 53–62, Amsterdam etc, 2002. IOS Press.

[4] T. Bench-Capon, K. Atkinson, and A. Chorley. Persuasion and value in legal argument. *Journal of Logic and Computation*, 15:1075–1097, 2005.

[5] T. Bench-Capon, T. Geldard, and P. Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8:233–254, 2000.

[6] T. Bench-Capon and H. Prakken. Argumentation. In A. Lodder and A. Oskamp, editors, *Information Technology and Lawyers: Advanced technology in the legal domain, from challenges to daily routine*, pages 61–80. Springer, Berlin, 2006.

[7] T. Bench-Capon and G. Sartor. A quantitative approach to theory coherence. In B. V. et al., editor, *Legal Knowledge and Information Systems. JURIX 2001: The Fourteenth Annual Conference*, pages 53–62, Amsterdam etc, 2001. IOS Press.

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

[9] T. Bench-Capon and G. Staniford. PLAID - proactive legal assistance. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 81–88, New York, 1995. ACM Press.

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

[11] F. Bex, S. v. d. Braak, H. v. Oostendorp, H. Prakken, B. Verheij, and G. Vreeswijk. Sense-making software for crime investigation: how to combine stories and arguments? *Law, Probability and Risk*, 6:145–168, 2007.

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

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

[14] P. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming, and n–person games. *Artificial Intelligence*, 77:321–357, 1995.

[15] A. Gardner. *Artificial Intelligence Approach to Legal Reasoning*. MIT Press, Cambridge, MA, 1987.

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

[17] T. Gordon. *The Pleadings Game. An Artificial Intelligence Model of Procedural Justice*. Kluwer Academic Publishers, Dordrecht/Boston/London, 1995.

[18] T. Gordon, H. Prakken, and D. Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171:875–896, 2007.

[19] J. Hage. A theory of legal reasoning and a logic to match. *Artificial Intelligence and Law*, 4:199–273, 1996.

[20] J. Hage. Formalizing legal coherence. In *Proceedings of the Eighth International Conference on Artificial Intelligence and Law*, pages 22–31, New York, 2001. ACM Press.

[21] R. Kowalski and F. Toni. Abstract argumentation. *Artificial Intelligence and Law*, 4:275–296, 1996.

[22] A. Lodder. *DiaLaw. On Legal Justification and Dialogical Models of Argumentation*. Kluwer Academic Publishers, Dordrecht/Boston/London, 1999.

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

[24] R. Loui, J. Norman, J. Alpeter, D. Pinkard, D. Craven, J. Linsday, and M. Foltz. Progress on Room 5: A testbed for public interactive semi-formal legal argumentation. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*, pages 207–214, New York, 1997. ACM Press.

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

[26] L. McCarty. An implementation of Eisner v. Macomber. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 276–286, New York, 1995. ACM Press.

[27] L. McCarty and N. Sridharan. The representation of an evolving system of legal concepts: II. Prototypes and deformations. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 246–253, 1981.

[28] L. Mommers. *Applied Legal Epistemology*. Doctoral dissertation Leiden University, 2002.

[29] A. Peczenik. Jumps and logic in the law. *Artificial Intelligence and Law*, 4:297–329, 1996.

[30] H. Prakken. A logical framework for modelling legal argument. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pages 1–9, New York, 1993. ACM Press.

[31] H. Prakken. Modelling defeasibility in law: logic or procedure? *Fundamenta Informaticae*, 48:253–271, 2001.

[32] H. Prakken. An exercise in formalising teleological case-based reasoning. *Artificial Intelligence and Law*, 10:113–133, 2002.

[33] H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*, pages 85–94, New York, 2005. ACM Press.

[34] H. Prakken. A formal model of adjudication dialogues. *Artificial Intelligence and Law*, 16, 2008. to appear.

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

[36] H. Prakken and G. Sartor. Formalising arguments about the burden of persuasion. In *Proceedings of the Eleventh International Conference on Artificial Intelligence and Law*, pages 97–106, New York, 2007. ACM Press.

[37] H. Prakken and G. Sartor. More on presumptions and burdens of proof. In G. Sartor, editor, *Legal Knowledge and Information Systems. JURIX 2008: The Twentyfirst Annual Conference*. IOS Press, Amsterdam etc., 2008.

[38] C. Reed and G. Rowe. Araucaria: Software for argument analysis, diagramming and representation. *International Journal of AI Tools*, 13:961–980, 2004.

[39] C. Reed, D. Walton, and F. Macagno. Argument diagramming in logic, law and artificial intelligence. *The Knowledge Engineering Review*, 22:87–109, 2007.

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

[41] R. Riveret, A. Rotolo, G. Sartor, H. Prakken, and B. Roth. Success chances in argument games: a probabilistic approach to legal disputes. In A. Lodder and L. Mommers, editors, *Legal Knowledge and Information Systems. JURIX 2007: The Twentieth Annual Conference*, pages 99–108. IOS Press, Amsterdam etc., 2007.

[42] D. Schum and P. Tillers. Marshaling evidence for adversary litigation. *Cardozo Law Review*, 13:657–704, 1991.

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

[44] S. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, 1958.

[45] B. Verheij. *Rules, reasons, arguments: formal studies of argumentation and defeat*. Doctoral dissertation University of Maastricht, 1996.

[46] B. Verheij. Automated argument assistance for lawyers. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, pages 43–52, New York, 1999. ACM Press.

[47] D. Walton. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, 1996.

[48] D. Walton and E. Krabbe. *Commitment in Dialogue. Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY, 1995.

[49] J. Wigmore. *The Principles of Judicial Proof*. Little, Brown and Company, Boston, 2nd edition, 1931.

[50] A. Wyner and T. Bench-Capon. Argument schemes for legal case-based reasoning. In A. Lodder and L. Mommers, editors, *Legal Knowledge and Information Systems. JURIX 2007: The Twentieth Annual Conference*, pages 139–149. IOS Press, Amsterdam etc., 2007.

[51] A. Wyner, T. Bench-Capon, and K. Atkinson. Arguments, values and baseballs: Representation of Popov v. Hayashi. In A. Lodder and L. Mommers, editors, *Legal Knowledge and Information Systems. JURIX 2007: The Twentieth Annual Conference*, pages 151–160. IOS Press, Amsterdam etc., 2007.

[52] J. Zeleznikow and A. Stranieri. The split-up system. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 185–195, New York, 1995. ACM Press.